



6400cl
5300xl
4200vl
3400cl

Management and Configuration Guide

ProCurve Switches

E.10.02 (Series 5300xl)

L.10.XX (Series 4200vl)

M.08.73 (Series 3400/6400cl)

www.procurve.com



ProCurve

Series 6400cl Switches

Series 5300xl Switches

Series 4200vl Switches

Series 3400cl Switches

October 2006

E.10.02 or Greater (5300xl)

L.10.01 or Greater (4200vl)

M.08.73 or Greater (3400/6400cl)

Management and Configuration Guide

© Copyright 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. All Rights Reserved.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard.

Publication Number

5990-6050
October 2006

Applicable Products

ProCurve Switch 5308XL	(J4819A)
ProCurve Switch 5372XL	(J4848A)
ProCurve Switch 5348XL	(J4849A)
ProCurve Switch 5304XL	(J4850A)
ProCurve Switch 3400cl-24G	(J4905A)
ProCurve Switch 3400cl-48G	(J4906A)
ProCurve Switch 4204vl	(J8770A)
ProCurve Switch 4208vl	(J8773A)
ProCurve Switch 4202vl-72	(J8772A)
ProCurve Switch 4202vl-48G	(J8771A)
ProCurve Switch 10G CX4 6400cl-6XG	(J8433A)
ProCurve Switch 10G X2 6400cl-6XG	(J8474A)

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are US registered trademarks of Microsoft Corporation. Internet Explorer is a trademark of Microsoft Corporation. Ethernet is a registered trademark of Xerox Corporation. Netscape is a registered trademark of Netscape Corporation. Cisco® is a trademark of Cisco Systems, Inc.

Disclaimer

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Contents

Product Documentation

About Your Switch Manual Set	xix
Feature Index	xx

1 Getting Started

Contents	1-1
Introduction	1-2
Conventions	1-2
Feature Descriptions by Model	1-2
Command Syntax Statements	1-3
Command Prompts	1-3
Screen Simulations	1-4
Port Identity Examples	1-4
Keys	1-4
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-9

2 Selecting a Management Interface

Contents	2-1
Overview	2-2
Understanding Management Interfaces	2-2
Advantages of Using the Menu Interface	2-3

Advantages of Using the CLI	2-4
General Benefits	2-4
Information on Using the CLI	2-5
Advantages of Using the Web Browser Interface	2-5
Advantages of Using ProCurve Manager or ProCurve Manager Plus	2-7

3 Using the Menu Interface

Contents	3-1
Overview	3-2
Starting and Ending a Menu Session	3-3
How To Start a Menu Interface Session	3-4
How To End a Menu Session and Exit from the Console	3-5
Main Menu Features	3-7
Screen Structure and Navigation	3-9
Rebooting the Switch	3-12
Menu Features List	3-14
Where To Go From Here	3-15

4 Using the Command Line Interface (CLI)

Contents	4-1
Overview	4-2
Accessing the CLI	4-2
Using the CLI	4-2
Privilege Levels at Logon	4-3
Privilege Level Operation	4-4
Operator Privileges	4-4
Manager Privileges	4-5
How To Move Between Levels	4-7
Listing Commands and Command Options	4-8
Listing Commands Available at Any Privilege Level	4-8
Listing Command Options	4-10

Displaying CLI “Help”	4-11
Configuration Commands and the Context Configuration Modes ..	4-12
Configuring Custom Login Banners for the Console and Web Browser Interfaces	4-15
Banner Operation with Telnet, Serial, or SSHv2 Access	4-16
Banner Operation with Web Browser Access	4-16
Configuring and Displaying a Non-Default Banner	4-16
Example of Configuring and Displaying a Banner	4-17
Operating Notes	4-19
CLI Control and Editing	4-21

5 Using the Web Browser Interface

Contents	5-1
Overview	5-2
General Features	5-3
Starting an Web Browser Interface Session with the Switch	5-4
Using a Standalone Web Browser in a PC or UNIX Workstation	5-4
Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)	5-5
Tasks for Your First Web Browser Interface Session	5-7
Viewing the “First Time Install” Window	5-7
Security: Creating Usernames and Passwords in the Browser Interface	5-8
Entering a User Name and Password	5-10
Using a User Name	5-10
If You Lose the Password	5-10
Online Help for the Web Browser Interface	5-11
Support/Mgmt URLs Feature	5-12
Support URL	5-13
Help and the Management Server URL	5-13
Using the PCM Server for Switch Web Help	5-14
Status Reporting Features	5-16
The Overview Window	5-16
The Port Utilization and Status Displays	5-17

Port Utilization	5-17
Port Status	5-19
The Alert Log	5-20
Sorting the Alert Log Entries	5-20
Alert Types and Detailed Views	5-21
The Status Bar	5-22
Setting Fault Detection Policy	5-24

6 Switch Memory and Configuration

Contents	6-1
Overview	6-3
Overview of Configuration File Management	6-3
Using the CLI To Implement Configuration Changes	6-6
Using the Menu and Web Browser Interfaces To Implement Configuration Changes	6-9
Menu: Implementing Configuration Changes	6-9
Using Save and Cancel in the Menu Interface	6-10
Rebooting from the Menu Interface	6-11
Web: Implementing Configuration Changes	6-12
Using Primary and Secondary Flash Image Options	6-13
Displaying the Current Flash Image Data	6-13
Switch Software Downloads	6-15
Local Switch Software Replacement and Removal	6-16
Rebooting the Switch	6-18
Operating Notes	6-21
Multiple Configuration Files on 5300xl and 4200vl Switches	6-22
General Operation	6-24
Transitioning to Multiple Configuration Files	6-26
Listing and Displaying Startup-Config Files	6-27
Viewing the Startup-Config File Status with Multiple Configuration Enabled	6-27
Displaying the Content of A Specific Startup-Config File	6-29
Changing or Overriding the Reboot Configuration Policy	6-29
Managing Startup-Config Files in the Switch	6-31

Renaming an Existing Startup-Config File	6-32
Creating a New Startup-Config File	6-32
Erasing a Startup-Config File	6-34
Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration	6-35
Transferring Startup-Config Files To or From a Remote Server	6-37
TFTP: Copying a Configuration File to a Remote Host	6-37
TFTP: Copying a Configuration File from a Remote Host	6-37
Xmodem: Copying a Configuration File to a Serially Connected Host	6-38
Xmodem: Copying a Configuration from a Serially Connected Host	6-38
Operating Notes for Multiple Configuration Files	6-39

7 Interface Access and System Information

Contents	7-1
Overview	7-2
Interface Access: Console/Serial Link, Web, and Inbound Telnet .	7-3
Menu: Modifying the Interface Access	7-4
CLI: Modifying the Interface Access	7-5
Denying Interface Access by Terminating Remote Management Sessions	7-8
System Information	7-9
Menu: Viewing and Configuring System Information	7-10
CLI: Viewing and Configuring System Information	7-11
Web: Configuring System Parameters	7-14

8 Configuring IP Addressing

Contents	8-1
Overview	8-2
IP Configuration	8-2
Just Want a Quick Start with IP Addressing?	8-3
IP Addressing with Multiple VLANs	8-4
Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL) ..	8-5

CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)	8-6
Web: Configuring IP Addressing	8-10
How IP Addressing Affects Switch Operation	8-11
DHCP/Bootp Operation	8-12
Network Preparations for Configuring DHCP/Bootp	8-14
IP Preserve: Retaining VLAN-1 IP Addressing Across	
Configuration File Downloads	8-15
Operating Rules for IP Preserve	8-16
Enabling IP Preserve	8-16

9 Time Protocols

Contents	9-1
Overview	9-2
TimeP Time Synchronization	9-2
SNTP Time Synchronization	9-3
Selecting a Time Synchronization Protocol or Turning Off Time	
Protocol Operation	9-3
General Steps for Running a Time Protocol on the Switch:	9-3
Disabling Time Synchronization	9-4
SNTP: Viewing, Selecting, and Configuring	9-5
Menu: Viewing and Configuring SNTP	9-6
CLI: Viewing and Configuring SNTP	9-8
Viewing the Current SNTP Configuration	9-9
Configuring (Enabling or Disabling) the SNTP Mode	9-10
TimeP: Viewing, Selecting, and Configuring	9-16
Menu: Viewing and Configuring TimeP	9-17
CLI: Viewing and Configuring TimeP	9-18
Viewing the Current TimeP Configuration	9-19
Configuring (Enabling or Disabling) the TimeP Mode	9-20
SNTP Unicast Time Polling with Multiple SNTP Servers	9-25
Address Prioritization	9-25
Displaying All SNTP Server Addresses Configured on the Switch . .	9-26
Adding and Deleting SNTP Server Addresses	9-26

Menu: Operation with Multiple SNTP Server Addresses Configured	9-28
SNTP Messages in the Event Log	9-28

10 Port Status and Basic Configuration

Contents	10-1
Overview	10-2
Viewing Port Status and Configuring Port Parameters	10-2
Menu: Port Configuration	10-6
CLI: Viewing Port Status and Configuring Port Parameters	10-8
Using the CLI To Enable or Disable Ports and Configure Port Mode	10-9
Enabling or Disabling Flow Control	10-11
Configuring a Broadcast Limit on the Switch	10-14
Configuring Auto-MDIX	10-15
Web: Viewing Port Status and Configuring Port Parameters	10-18
Using Friendly (Optional) Port Names	10-18
Configuring and Operating Rules for Friendly Port Names	10-19
Configuring Friendly Port Names	10-19
Displaying Friendly Port Names with Other Port Data	10-21

11 Power Over Ethernet (PoE) Operation for the Series 5300xl Switches

Contents	11-1
PoE Operation on the Series 5300xl Switches	11-2
Introduction	11-2
PoE Terminology	11-3
Overview of Operation	11-4
Related Publications	11-4
General PoE Operation	11-5
Configuration Options	11-5
PD Support	11-6
Power Priority Operation	11-8

Configuring PoE Operation	11-10
Changing the PoE Port Priority Level	11-10
Disabling or Re-Enabling PoE Port Operation	11-11
Changing the Threshold for Generating a Power Notice	11-11
Configuring Optional PoE Port Identifiers	11-12
Viewing PoE Configuration and Status	11-15
Displaying the Switch's Global PoE Power Status	11-15
Displaying an Overview of PoE Status on All Ports	11-16
Displaying the PoE Status on Specific Ports	11-17
Planning and Implementing a PoE Configuration	11-19
Assigning PoE Ports to VLANs	11-19
Applying Security Features to PoE Configurations	11-20
Assigning Priority Policies to PoE Traffic	11-20
Calculating the Maximum Load for an xl PoE Module	11-21
PoE Operating Notes	11-23
PoE Event Log Messages	11-23
"Informational" PoE Event-Log Messages	11-23
"Warning" PoE Event-Log Messages	11-25

12 Access Controller xl Module for the Series 5300xl Switches

Contents	12-1
Introduction	12-3
General Operation	12-3
Related Publications	12-3
Terminology	12-4
Access Controller xl Module Overview	12-5
Module Operation	12-5
Using 5300xl Features with the Access Controller xl Module ...	12-7
Routing Infrastructure Support	12-10
Using 5300xl Switch Network Address Translation with the ACM	12-11
The Role of VLANs	12-12
Client VLANs	12-12
Static VLAN Features Supported on Client VLANs	12-13

General Operating Rules	12-14
Configuring the ACM on the Network	12-14
Configuring the Access Controller x1 Module	12-16
Configuring Downlink Client Ports	12-16
Changing the VLAN-Base	12-18
Configuring Client VLANs	12-19
Configuring Uplink Network Ports	12-19
Configuring the Uplink VLAN	12-19
ACM Configuration Commands Summary and Syntax	12-20
Configuration Context Command Syntax	12-20
Access Controller Context Command Syntax	12-22
Displaying Access Controller x1 Status from the 5300x1 CLI ...	12-24
ACM Display Commands Summary and Syntax	12-24
Configuration Context Command Syntax	12-25
Access Controller Context Command Syntax	12-26
Managing the ACM	12-27
Using the ACM's Extended CLI	12-27
Downloading New Software to the Module	12-30
Resetting the Module to Factory Defaults	12-30
Operating Notes	12-31
BIOS POST Event Log Messages	12-32

13 Port Trunking

Contents	13-1
Overview	13-2
Port Trunk Features and Operation	13-4
Trunk Configuration Methods	13-5
Menu: Viewing and Configuring a Static Trunk Group	13-9
CLI: Viewing and Configuring Port Trunk Groups	13-11
Using the CLI To View Port Trunks	13-11
Using the CLI To Configure a Static or Dynamic Trunk Group ...	13-14
Web: Viewing Existing Port Trunk Groups	13-17

Trunk Group Operation Using LACP	13-18
Default Port Operation	13-20
LACP Notes and Restrictions	13-21
Trunk Group Operation Using the “Trunk” Option	13-24
How the Switch Lists Trunk Data	13-25
Outbound Traffic Distribution Across Trunked Links	13-25

14 Port Traffic Controls

Contents	14-1
Overview	14-3
All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches	14-4
Introduction	14-4
Rate-Limiting Operation	14-5
Configuring Inbound Rate-Limiting	14-5
Displaying the Current Rate-Limit Configuration	14-6
Operating Notes for Rate-Limiting	14-7
ICMP Rate-Limiting	14-10
Terminology	14-10
Effect of ICMP Rate-Limiting	14-10
Operating Notes for ICMP Rate-Limiting	14-17
Guaranteed Minimum Bandwidth (GMB) on the Series 5300xl Switches	14-21
Introduction	14-21
Terminology	14-21
GMB Operation	14-21
Configuring Guaranteed Minimum Bandwidth for Outbound Traffic	14-23
Displaying the Current Guaranteed Minimum Bandwidth Configuration	14-25
GMB Operating Notes	14-26
Jumbo Packets on the Series 3400cl and Series 6400cl Switches	14-27
Terminology	14-27
Operating Rules	14-28

Configuring Jumbo Packet Operation	14-29
Overview	14-29
Viewing the Current Jumbo Configuration	14-29
Enabling or Disabling Jumbo Traffic on a VLAN	14-31
Operating Notes for Jumbo Traffic-Handling	14-32
Troubleshooting	14-34

15 Configuring for Network Management Applications

Contents	15-1
Using SNMP Tools To Manage the Switch	15-3
Overview	15-3
SNMP Management Features	15-4
Configuring for SNMP Access to the Switch	15-4
Configuring for SNMP Version 3 Access to the Switch	15-5
SNMP Version 3 Commands	15-6
Enabling SNMPv3	15-7
SNMPv3 Users	15-7
Group Access Levels	15-10
SNMPv3 Communities	15-11
Menu: Viewing and Configuring non-SNMP version 3 Communities	15-13
CLI: Viewing and Configuring SNMP Community Names	15-14
SNMPv3 Notification and Traps	15-16
SNMPv1 and SNMPv2c Trap Features	15-19
CLI: Configuring and Displaying Trap Receivers	15-19
Using the CLI To Enable Authentication Traps	15-22
Advanced Management: RMON	15-22
LLDP (Link-Layer Discovery Protocol)	15-24
Terminology	15-25
General LLDP Operation	15-27
LLDP-MED	15-27
Packet Boundaries in a Network Topology	15-27
Configuration Options	15-28
Options for Reading LLDP Information Collected by the Switch ..	15-30
LLDP and LLDP-MED Standards Compatibility	15-31

LLDP Operating Rules	15-31
LLDP Data Management on the Series 3400cl and 6400cl Switches	15-32
LLDP Neighbor Data	15-32
Configuring LLDP Operation	15-33
Viewing the Current Configuration	15-34
Configuring Global LLDP Packet Controls	15-37
Configuring SNMP Notification Support	15-41
Changing the Minimum Interval for Successive Data Change Notifications for the Same Neighbor	15-42
Configuring Per-Port Transmit and Receive Modes	15-42
Configuring Basic LLDP Per-Port Advertisement Content	15-43
Configuring Support for Port Speed and Duplex Advertisements on the 5300xl and 4200vl Switches	15-46
LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches	15-47
LLDP-MED Topology Change Notification	15-50
LLDP-MED Fast Start Control	15-51
Advertising Device Capability, Network Policy, PoE Status and Location Data	15-52
Configuring Location Data for LLDP-MED Devices	15-56
Displaying Advertisement Data	15-62
Displaying Switch Information Available for Outbound Advertisements	15-63
Displaying LLDP Statistics	15-68
LLDP Operating Notes	15-70
LLDP and CDP Data Management	15-72
LLDP and CDP Neighbor Data	15-72
CDP Operation and Commands	15-74

A File Transfers

Contents	A-1
Overview	A-3
Downloading Switch Software	A-3
General Software Download Rules	A-4
Using TFTP To Download Switch Software from a Server	A-4

Menu: TFTP Download from a Server to Primary Flash	A-5
CLI: TFTP Download from a Server to Flash	A-6
Using Secure Copy and SFTP	A-8
How It Works	A-9
The SCP/SFTP Process	A-10
Disable TFTP and Auto-TFTP for Enhanced Security	A-10
Command Options	A-13
Authentication	A-14
SCP/SFTP Operating Notes	A-14
Using Xmodem to Download Switch Software From a PC or UNIX Workstation	A-16
Menu: Xmodem Download to Primary Flash	A-16
CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash	A-17
Switch-to-Switch Download	A-18
Menu: Switch-to-Switch Download to Primary Flash	A-18
CLI: Switch-To-Switch Downloads	A-19
Using PCM+ to Update Switch Software	A-21
Troubleshooting TFTP Downloads	A-21
Transferring Switch Configurations and ACL Command Files . .	A-23
TFTP: Copying a Configuration from a Remote Host	A-23
TFTP: Copying a Configuration File to a Remote Host	A-24
TFTP: Uploading an ACL Command File from a TFTP Server	A-24
Xmodem: Copying a Configuration File from the Switch to a Serially Connected PC or UNIX Workstation	A-27
Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation	A-27
Copying Diagnostic Data to a Remote Host, PC, or UNIX Workstation	A-29
Copying Command Output to a Destination Device	A-29
Copying Event Log Output to a Destination Device	A-30
Copying Crash Data Content to a Destination Device	A-30
Copying Crash Log Data Content to a Destination Device	A-31

B Monitoring and Analyzing Switch Operation

Contents	B-1
Overview	B-3
Status and Counters Data	B-4
Menu Access To Status and Counters	B-5
General System Information	B-5
Menu Access	B-5
CLI Access	B-6
Switch Management Address Information	B-6
Menu Access	B-6
CLI Access	B-7
Module Information	B-8
Menu: Displaying Port Status	B-8
CLI Access	B-8
Port Status	B-9
Menu: Displaying Port Status	B-9
CLI Access	B-9
Web Access	B-9
Viewing Port and Trunk Group Statistics and Flow Control Status	B-10
Menu Access to Port and Trunk Statistics	B-11
CLI Access To Port and Trunk Group Statistics	B-12
Web Browser Access To View Port and Trunk Group Statistics	B-12
Viewing the Switch's MAC Address Tables	B-13
Menu Access to the MAC Address Views and Searches	B-13
CLI Access for MAC Address Views and Searches	B-16
Spanning Tree Protocol (STP) Information	B-17
Menu Access to STP Data	B-17
CLI Access to STP Data	B-18
Internet Group Management Protocol (IGMP) Status	B-19
VLAN Information	B-20
Web Browser Interface Status Information	B-22

Interface Monitoring Features	B-23
Menu: Configuring Port and Static Trunk Monitoring	B-24
CLI: Configuring Port, Mesh, and Static Trunk Monitoring	B-26
Web: Configuring Port Monitoring	B-29

C Troubleshooting

Contents	C-1
Overview	C-3
Troubleshooting Approaches	C-4
Browser or Telnet Access Problems	C-5
Unusual Network Activity	C-7
General Problems	C-7
802.1Q Prioritization Problems	C-8
ACL Problems	C-8
IGMP-Related Problems	C-13
LACP-Related Problems	C-13
Mesh-Related Problems	C-14
Port-Based Access Control (802.1x)-Related Problems	C-15
QoS-Related Problems	C-18
Radius-Related Problems	C-18
Spanning-Tree Protocol (STP) and Fast-Uplink Problems	C-19
SSH-Related Problems	C-20
TACACS-Related Problems	C-22
TimeP, SNTP, or Gateway Problems	C-24
VLAN-Related Problems	C-24
Using the Event Log To Identify Problem Sources	C-27
Menu: Entering and Navigating in the Event Log	C-29
CLI: Listing Events	C-30
Reducing Duplicate Event Log and SNMP Trap Messages	C-31
Debug and Syslog Messaging Operation	C-34
Debug Command Operation	C-35
Debug Types	C-36
Debug Destinations	C-38
Syslog Operation	C-39

Viewing the Debug Configuration	C-40
Steps for Configuring Debug and Syslog Messaging	C-40
Operating Notes for Debug and Syslog	C-44
Diagnostic Tools	C-45
Port Auto-Negotiation	C-45
Ping and Link Tests	C-45
Web: Executing Ping or Link Tests	C-47
CLI: Ping or Link Tests	C-48
Displaying the Configuration File	C-50
CLI: Viewing the Configuration File	C-50
Web: Viewing the Configuration File	C-50
Listing Switch Configuration and Operation Details	C-50
CLI Administrative and Troubleshooting Commands	C-52
Traceroute Command	C-53
Restoring the Factory-Default Configuration	C-57
CLI: Resetting to the Factory-Default Configuration	C-57
Clear/Reset: Resetting to the Factory-Default Configuration ..	C-57
Restoring a Flash Image	C-58

D MAC Address Management

Contents	D-1
Overview	D-2
Determining MAC Addresses	D-3
Menu: Viewing the Switch's MAC Addresses	D-4
CLI: Viewing the Port and VLAN MAC Addresses	D-5
Viewing the MAC Addresses of Connected Devices	D-8

E Daylight Savings Time on ProCurve Switches

Index

Product Documentation

About Your Switch Manual Set

Note

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features, please visit the ProCurve Networking Web site at www.procurve.com, click on **Technical support**, and then click on **Product manuals (all)**.

Printed Publications

The two publications listed below are printed and shipped with your switch. The latest version of each is also available in PDF format on the ProCurve Networking Web sit, as described in the above Note.

- *Read Me First*—Provides software update information, product notes, and other information.
- *Installation and Getting Started Guide*—Explains how to prepare for and perform the physical installation and connect the switch to your network.

Electronic Publications

The latest version of each of the publications listed below is available in PDF format on the ProCurve Web site, as described in the above Note.

- *Management and Configuration Guide*—Describes how to configure, manage, and monitor switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features such as STP, QoS, and IP routing.
- *Access Security Guide*—Explains how to configure access security features and user authentication on the switch.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the main product guide.

Feature Index

For the manual set supporting your switch model, the following feature index indicates which manual to consult for information on a given software feature and which switches support that feature.

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide	Supported on 5300xl	Supported on 4200vl	Supported on 3400cl/ 6400cl
802.1Q VLAN Tagging		X		yes	yes	yes
802.1X Port-Based Priority	X			yes	yes	yes
802.1X Multiple Authenticated Clients per port			X	yes	no	yes
ACLs		X		yes	no	yes
AAA Authentication			X	yes	yes	yes
Authorized IP Managers			X	yes	yes	yes
Authorized Manager List (web, telnet, TFTP)			X	yes	yes	yes
Auto MDIX Configuration	X			yes	yes	yes
BOOTP	X			yes	yes	yes
Config File	X			yes	yes	yes
Console Access	X			yes	yes	yes
Copy Command	X			yes	yes	yes
CoS (Class of Service)		X		yes	yes	yes
Debug	X			yes	yes	yes
DHCP Configuration		X		yes	yes	yes
DHCP Option 82		X		yes	yes	no
DHCP/Bootp Operation	X			yes	yes	yes
Diagnostic Tools	X			yes	yes	yes
Downloading Software	X			yes	yes	yes

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide	Supported on 5300xl	Supported on 4200vl	Supported on 3400cl/ 6400cl
Eavesdrop Protection			X	yes	yes	no
Event Log	X			yes	yes	yes
Factory Default Settings	X			yes	yes	yes
Flow Control (802.3x)	X			yes	yes	yes
File Management	X			yes	yes	yes
File Transfers	X			yes	yes	yes
Friendly Port Names	X			yes	yes	yes
Guaranteed Minimum Bandwidth (GMB)	X			yes	no	yes
GVRP		X		yes	yes	yes
IGMP		X		yes	yes	yes
Delayed Group Flush		X		yes	yes	yes
Interface Access (Telnet, Console/ Serial, Web)	X			yes	yes	yes
IP Addressing	X			yes	yes	yes
IP Routing		X		yes	yes	yes
Jumbo Support		X		no	no	yes
LACP	X			yes	yes	yes
Link	X			yes	yes	yes
LLDP	X			yes	yes	yes
LLDP-Med	X			yes	yes	no
MAC Address Management	X			yes	yes	yes
MAC Lockdown			X	yes	yes	yes
MAC Lockout			X	yes	yes	yes
MAC-based Authentication			X	yes	yes	yes
MAC authentication RADIUS support			X	yes	yes	yes
Management VLAN		X		yes	yes	yes

Product Documentation

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide	Supported on 5300xl	Supported on 4200vl	Supported on 3400cl/ 6400cl
Meshing		X		yes	no	yes
Monitoring and Analysis	X			yes	yes	yes
Multicast Filtering			X	yes	no	no
Multiple Configuration Files	X			yes	yes	yes
Network Management Applications	X			yes	SNMP only	SNMP only
OpenView Device Management	X			yes	yes	yes
OSPF		X		yes	no	yes
Passwords			X	yes	yes	yes
Password Clear Protection			X	yes	yes	yes
PCM	X			yes	yes	yes
PIM Dense		X		yes	no	no
Ping	X			yes	yes	yes
Port Configuration	X			yes	yes	yes
Port Monitoring		X		yes	yes	yes
Port Security			X	yes	yes	yes
Port Status	X			yes	yes	yes
Port Trunking (LACP)	X			yes	yes	yes
Port-Based Access Control			X	yes	yes	yes
Port-Based Priority (802.1Q)	X			yes	yes	yes
Power over Ethernet (PoE)	X			yes	no	no
Protocol Filters			X	yes	no	no
Protocol VLANs		X		yes	no	yes
Quality of Service (QoS)		X		yes	yes	yes
RADIUS Authentication and Accounting			X	yes	yes	yes
Rate-limiting	X			yes	no	yes
RIP		X		yes	no	yes

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide	Supported on 5300xl	Supported on 4200vl	Supported on 3400cl/ 6400cl
RMON 1,2,3,9	X			yes	yes	yes
Routing		X		yes	yes	yes
Routing - IP Static		X		yes	yes	yes
Secure Copy	X			yes	yes	yes
SFLOW				yes	yes	yes
SFTP	X			yes	yes	yes
SNMPv3	X			yes	yes	yes
Software Downloads (SCP/SFTP, TFTP, Xmodem)	X			yes	yes	yes
Source-Port Filters			X	yes	yes	yes
Spanning Tree (STP, RSTP, MSTP)		X		yes	yes	yes
SSHv2 (Secure Shell) Encryption			X	yes	yes	yes
SSL (Secure Socket Layer)			X	yes	yes	yes
Stack Management (Stacking)		X		no	yes	yes
Syslog	X			yes	yes	yes
System Information	X			yes	yes	yes
TACACS+ Authentication			X	yes	yes	yes
Telnet Access	X			yes	yes	yes
TFTP	X			yes	yes	yes
Time Protocols (TimeP, SNTP)	X			yes	yes	yes
Traffic/Security Filters			X	yes	yes	yes
Troubleshooting	X			yes	yes	yes
UDP Forwarder		X		yes	yes	yes
Virtual Stacking		X		no	yes	yes
Virus Throttling (connection-rate filtering)			X	yes	no	no
VLANs		X		yes	yes	yes
VLAN Mirroring (1 static VLAN)		X		yes	yes	no

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide	Supported on 5300xl	Supported on 4200vl	Supported on 3400cl/ 6400cl
Voice VLAN		X		yes	yes	yes
Web Authentication RADIUS Support			X	yes	yes	yes
Web-based Authentication			X	yes	yes	yes
Web UI	X			yes	yes	yes
Xmodem	X			yes	yes	yes
XRRP		X		yes	no	yes

Getting Started

Contents

Introduction	1-2
Conventions	1-2
Feature Descriptions by Model	1-2
Command Syntax Statements	1-3
Command Prompts	1-3
Screen Simulations	1-4
Port Identity Examples	1-4
Keys	1-4
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-9

Introduction

This *Management and Configuration Guide* is intended for use with the following switches:

- | | |
|---|--|
| ■ ProCurve Switch 10G CX4
6400cl-6xg | ■ ProCurve Switch 10G X2
6400cl-6xg |
| ■ ProCurve Switch 5304xl | ■ ProCurve Switch 5348xl |
| ■ ProCurve Switch 5308xl | ■ ProCurve Switch 5372xl |
| ■ ProCurve Switch 4204vl | ■ ProCurve Switch 4208vl |
| ■ ProCurve Switch 4202vl-48G | ■ ProCurve Switch 4202vl-72 |
| ■ ProCurve Switch 3400cl-24G | ■ ProCurve Switch 3400cl-48G |

This guide describes how to use the command line interface (CLI), Menu interface, and web browser to configure, manage, monitor, and troubleshoot switch operation.

For an overview of other product documentation for the above switches, refer to “*Product Documentation*” on page xix.

The *Product Documentation CD-ROM* shipped with the switch includes a copy of this guide. You can also download a copy from the ProCurve Networking web site, **www.procurve.com**.

Conventions

This guide uses the following conventions for command syntax and displayed information.

Feature Descriptions by Model

In cases where a software feature is not available in all of the switch models covered by this guide, the section heading specifically indicates which product or product series offer the feature.

For example, (the switch is highlighted here in **bold italics**):

“QoS Pass-Through Mode on the **Series 5300xl and 4200vl Switches**”.

Command Syntax Statements

Syntax: ip default-gateway < ip-addr >

Syntax: show interfaces [*port-list*]

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate optional elements.
- Braces (< >) enclose required elements.
- Braces within square brackets ([< >]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:
 “Use the **copy tftp** command to download the key from a TFTP server.”
- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, you must provide one or more port numbers:

Syntax: aaa port-access authenticator < *port-list* >

Command Prompts

In the default configuration, your switch displays one of the following CLI prompts:

```
ProCurve 6400cl#
ProCurve 5304xl#
ProCurve 5308xl#
ProCurve 4204vl#
ProCurve 4208vl#
ProCurve 3400-24cl#
ProCurve 3400-48cl#
```

To simplify recognition, this guide uses **ProCurve** to represent command prompts for all models. For example:

```
ProCurve#
```

(You can use the **hostname** command to change the text in the CLI prompt.)

Screen Simulations

Displayed Text. Figures containing simulated screen text and command output look like this:

```
ProCurve> show version
Image stamp:    /sw/code/build/info
                September 30 2005 13:43:13
                E.08.22
                139

ProCurve>
```

Figure 1-1. Example of a Figure Showing a Simulated Screen

In some cases, brief command-output sequences appear without figure identification. For example:

```
ProCurve(config)# clear client-public-key
ProCurve(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Port Identity Examples

This guide describes software applicable to both chassis-based and stackable ProCurve switches. Where port identities are needed in an example, this guide uses the chassis-based port identity system, such as “A1”, “B3-B5”, “C7”, etc. However, unless otherwise noted, such examples apply equally to the stackable switches, which typically use only numbers, such as “1”, “3-5”, “15”, etc. for port identities.

Keys

Simulations of actual keys use a bold, sans-serif typeface with square brackets. For example, the Tab key appears as **[Tab]** and the “Y” key appears as **[Y]**.

Sources for More Information

For additional information about switch operation and features not covered in this guide, consult the following sources:

- For information on which product manual to consult on a given software feature, refer to the chapter “*Product Documentation*”.

Note

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features, visit the ProCurve Networking web site at www.procurve.com, click on **Technical support**, and then click on **Product Manuals (all)**.

- Software Release Notes—Release notes are posted on the ProCurve Networking web site and provide information on new software updates:
 - New features and how to configure and use them
 - Software management, including downloading software to the switch
 - Software fixes addressed in current and previous releases

To view and download a copy of the latest software release notes for your switch, refer to “Getting Documentation From the Web” on page 1-6.

- Product Notes and Software Update Information—The printed *Read Me First* shipped with your switch provides software update information, product notes, and other information. For the latest version, refer to “Getting Documentation From the Web” on page 1-6.
- Installation and Getting Started Guide—Use the *Installation and Getting Started Guide* shipped with your switch to prepare for and perform the physical installation. This guide also steps you through connecting the switch to your network and assigning IP addressing, as well as describing the LED indications for correct operation and trouble analysis. A PDF version of this guide is also provided on the *Product Documentation CD-ROM* shipped with the switch. And you can download a copy from the ProCurve Networking web site. (See “Getting Documentation From the Web” on page 1-6.)
- Advanced Traffic Management Guide—Use the *Advanced Traffic Management Guide* for information on:
 - VLANs: Static port-based and protocol VLANs, and dynamic GVRP VLANs
 - Multicast traffic control (IGMP) and Protocol-Independent Multicast routing (PIM-DM)
 - Spanning-Tree: 802.1D (STP), 802.1w (RSTP), and 802.1s (MSTP)
 - Meshing
 - Quality-of-Service (QoS)
 - Access Control Lists (ACLs)
 - IP routing
 - Static NAT for intranet applications (*Series 5300xl switches only*)
 - XRRP (XL Router Redundancy Protocol)

- Access Security Guide—Use the *Access Security Guide* for information on:
 - Local username and password security
 - Web-Based and MAC-based authentication
 - RADIUS and TACACS+ authentication
 - SSH (Secure Shell) and SSL (Secure Socket Layer) operation
 - 802.1x port-based access control
 - Port security operation with MAC-based control
 - Authorized IP Manager security
 - Key Management System (KMS)

Getting Documentation From the Web

1. Go to the ProCurve Networking web site at
www.procurve.com
2. Click on **Technical support**.
3. Click on **Product manuals (all)**.
4. Click on the product for which you want to view or download a manual.

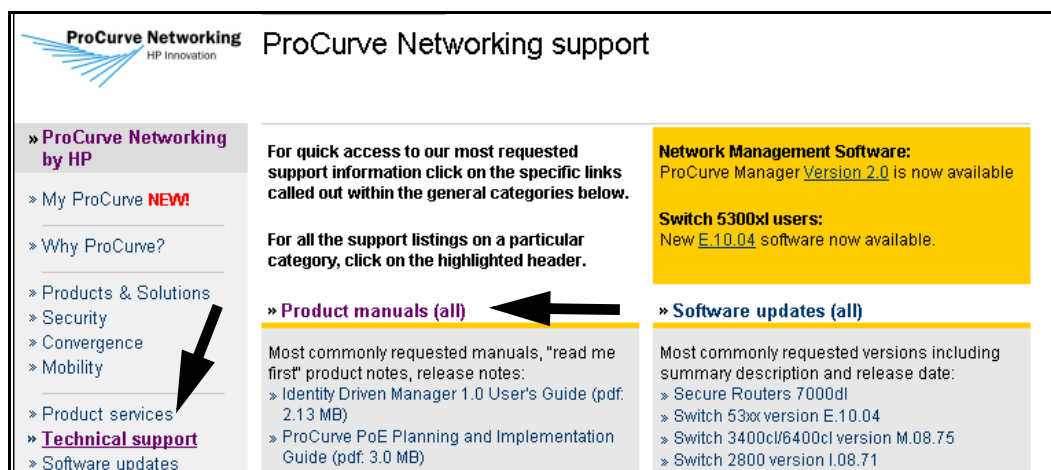


Figure 1-2. Example of How to Locate Product Manuals on the ProCurve Networking Web Site

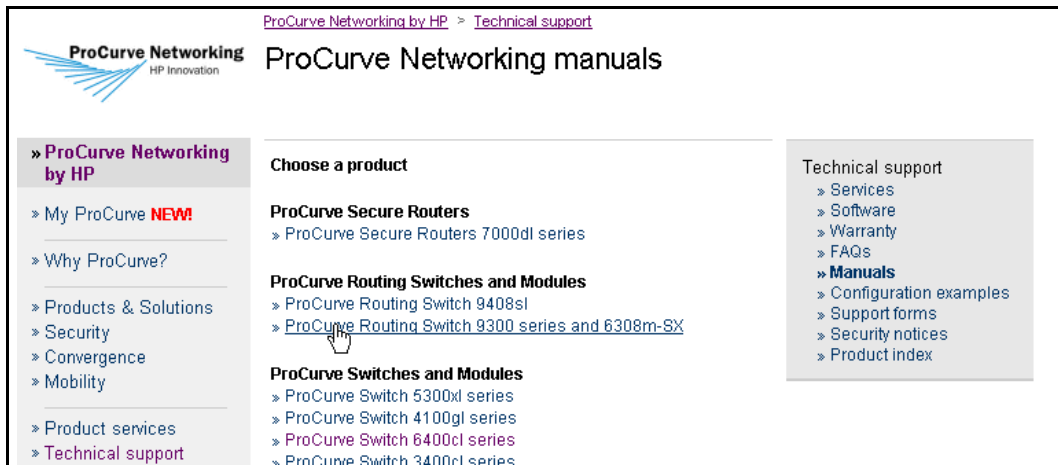
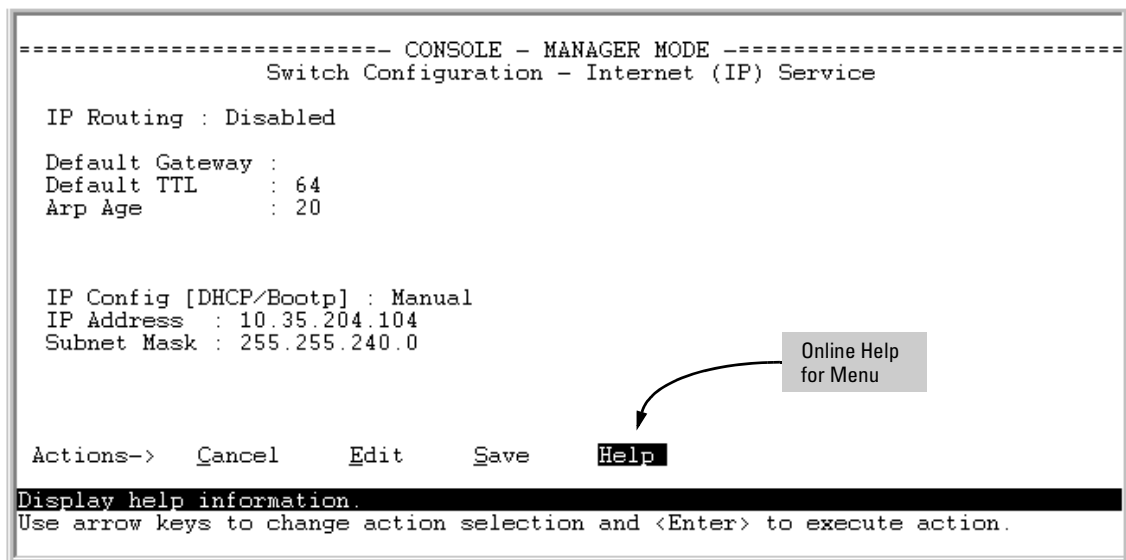


Figure 1-3. Listing of ProCurve Manuals on the ProCurve Networking Web Site

Online Help

If you need information on specific parameters in the menu interface, refer to the online help provided in the interface. For example:



If you need information on a specific command in the CLI, type the command name followed by “help”. For example:

```
ProCurve# write help
Usage: write <memory|terminal>

Description: View or save the running configuration of the switch.

        write terminal - displays the running configuration of the
                        switch on the terminal
        write memory   - saves the running configuration of the
                        switch to flash. The saved configuration
                        becomes the boot-up configuration of the switch
                        the next time it is booted.
```

If you need information on specific features in the web browser interface, use the online help available for the web browser interface. For more information on web browser Help options, refer to “Online Help for the Web Browser Interface” on page 5-11.

If you need further information on ProCurve switch technology, visit the ProCurve Networking web site at:

www.procurve.com

Need Only a Quick Start?

IP Addressing

If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using VLANs, ProCurve recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter **setup** at the CLI Manager level prompt.
Procurve# setup
- In the Main Menu of the Menu interface, select

8. Run Setup

For more on using the Switch Setup screen, see the *Installation and Getting Started Guide* you received with the switch.

To Set Up and Install the Switch in Your Network

Use the ProCurve *Installation and Getting Started Guide* (shipped with the switch) for the following:

- Notes, cautions, and warnings related to installing and using the switch and its related modules
- Instructions for physically installing the switch in your network
- Quickly assigning an IP address and subnet mask, set a Manager password, and (optionally) configure other basic features.
- Interpreting LED behavior.

For the latest version of the *Installation and Getting Started Guide* for your switch, refer to “Getting Documentation From the Web” on page 1-6.

Getting Started

To Set Up and Install the Switch in Your Network

— *This page unused intentionally*—

Selecting a Management Interface

Contents

Overview	2-2
Understanding Management Interfaces	2-2
Advantages of Using the Menu Interface	2-3
Advantages of Using the CLI	2-4
General Benefits	2-4
Information on Using the CLI	2-5
Advantages of Using the Web Browser Interface	2-5
Advantages of Using ProCurve Manager or ProCurve Manager Plus	2-7

Overview

This chapter describes the following:

- Management interfaces for the switches covered by this guide
 - Advantages of using each interface
-

Understanding Management Interfaces

Management interfaces enable you to reconfigure the switch and to monitor switch status and performance. The switch offers the following interfaces:

- **Menu interface**—a menu-driven interface offering a subset of switch commands through the built-in VT-100/ANSI console—**2-3**
- **CLI**—a command line interface offering the full set of switch commands through the VT-100/ANSI console built into the switch—**2-4**
- **Web browser interface**—a switch interface offering status information and a subset of switch commands through a standard web browser (such as Netscape Navigator or Microsoft Internet Explorer)—**2-5**
- **ProCurve Manager (PCM)**—a windows-based network management solution included in-box with all manageable ProCurve devices. Features include automatic device discovery, network status summary, topology and mapping, and device management.
- **ProCurve Manager Plus (PCM+)**—a complete windows-based network management solution that provides both the basic features offered with PCM, as well as more advanced management features, including in-depth traffic analysis, group and policy management, configuration management, device software updates, and advanced VLAN management. (ProCurve includes a copy of PCM+ in-box for a free 30-day trial.)

This manual describes how to use the menu interface (chapter 2), the CLI (chapter 3), the web browser interface (chapter 4), and how to use these interfaces to configure and monitor the switch.

For information on how to access the web browser interface Help, see “Online Help for the Web Browser Interface” on page 5-11.

To use ProCurve Manager or ProCurve Manager Plus, refer to the *Getting Started Guide* and the *Administrator's Guide*, which are available electronically with the software for these applications. For more information, visit the ProCurve Networking web site at www.procurve.com.

Advantages of Using the Menu Interface

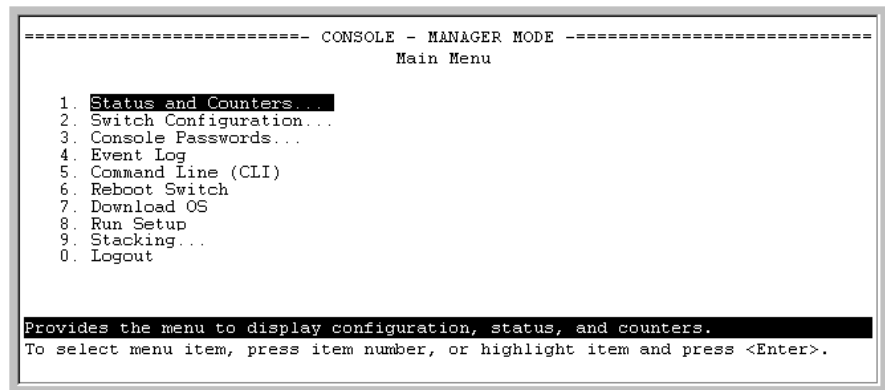


Figure 2-1. Example of the Console Interface Display (3400cl/6400cl Switches)

- **Provides quick, easy management access** to a menu-driven subset of switch configuration and performance features:

- IP addressing
- VLANs and GVRP
- Port Security
- Port and Static Trunk Group
- Spanning Tree
- System information
- Local passwords
- SNMP communities
- Time protocols
- Stacking (3400cl/6400cl/4200v1 switches only)

The menu interface also provides access for:

- Setup screen
- Event Log display
- Switch and port status displays
- Switch and port statistic and counter displays
- Reboots
- Software downloads

- **Offers out-of-band access** (through the RS-232 connection) to the switch, so network bottlenecks, crashes, lack of configured or correct IP address, and network downtime do not slow or prevent access
- **Enables Telnet (in-band) access** to the menu functionality.
- **Allows faster navigation**, avoiding delays that occur with slower display of graphical objects over a web browser interface.
- **Provides more security**; configuration information and passwords are not seen on the network.

Advantages of Using the CLI

ProCurve>	Prompt for Operator Level
ProCurve#	Prompt for Manager Level
ProCurve(config) #	Prompt for Global Configuration Level
ProCurve(<context>) #	Prompt for Context Configuration Levels
For example:	
ProCurve(eth-1-5) #	
ProCurve(vlan-1) #	
ProCurve(pim) #	
ProCurve(rip) #	

Figure 2-2. Command Prompt Examples

General Benefits

- Provides access to the complete set of the switch configuration, performance, and diagnostic features.
- Offers out-of-band access (through the RS-232 connection) or Telnet (in-band) access.
- Enables quick, detailed system configuration and management access to system operators and administrators experienced in command prompt interfaces.
- Provides help at each level for determining available options and variables.

Information on Using the CLI

- For information on how to use the CLI, refer to chapter 3. “Using the Command Line Interface (CLI)”.
- To perform specific procedures (such as configuring IP addressing or VLANs), use the Contents listing at the front of the manual to locate the information you need.
- For monitoring and analyzing switch operation, refer to appendix B.
- For information on individual CLI commands, refer to the Index or to the online Help provided in the CLI interface.

Advantages of Using the Web Browser Interface

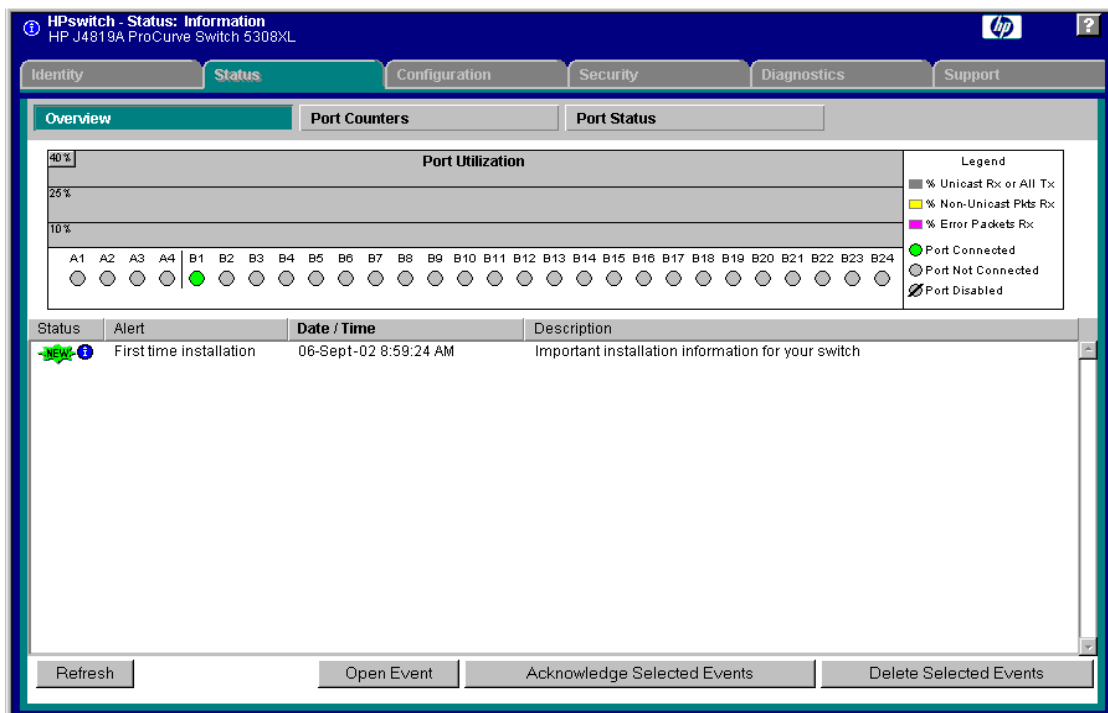


Figure 2-3. Example of the Web Browser Interface

- **Easy access** to the switch from anywhere on the network

Selecting a Management Interface

Advantages of Using the Web Browser Interface

- **Familiar browser interface**—locations of window objects consistent with commonly used browsers, uses mouse clicking for navigation, no terminal setup
- **Many features have all their fields in one screen** so you can view all values at once
- **More visual cues**, using colors, status bars, device icons, and other graphical objects instead of relying solely on alphanumeric values
- **Display of acceptable ranges of values available** in configuration list boxes

Advantages of Using ProCurve Manager or ProCurve Manager Plus

You can operate ProCurve Manager and ProCurve Manager Plus (PCM and PCM+) from a PC on the network to monitor traffic, manage your hubs and switches, and proactively recommend network changes to increase network uptime and optimize performance. Easy to install and use, PCM and PCM+ are the answers to your management challenges.

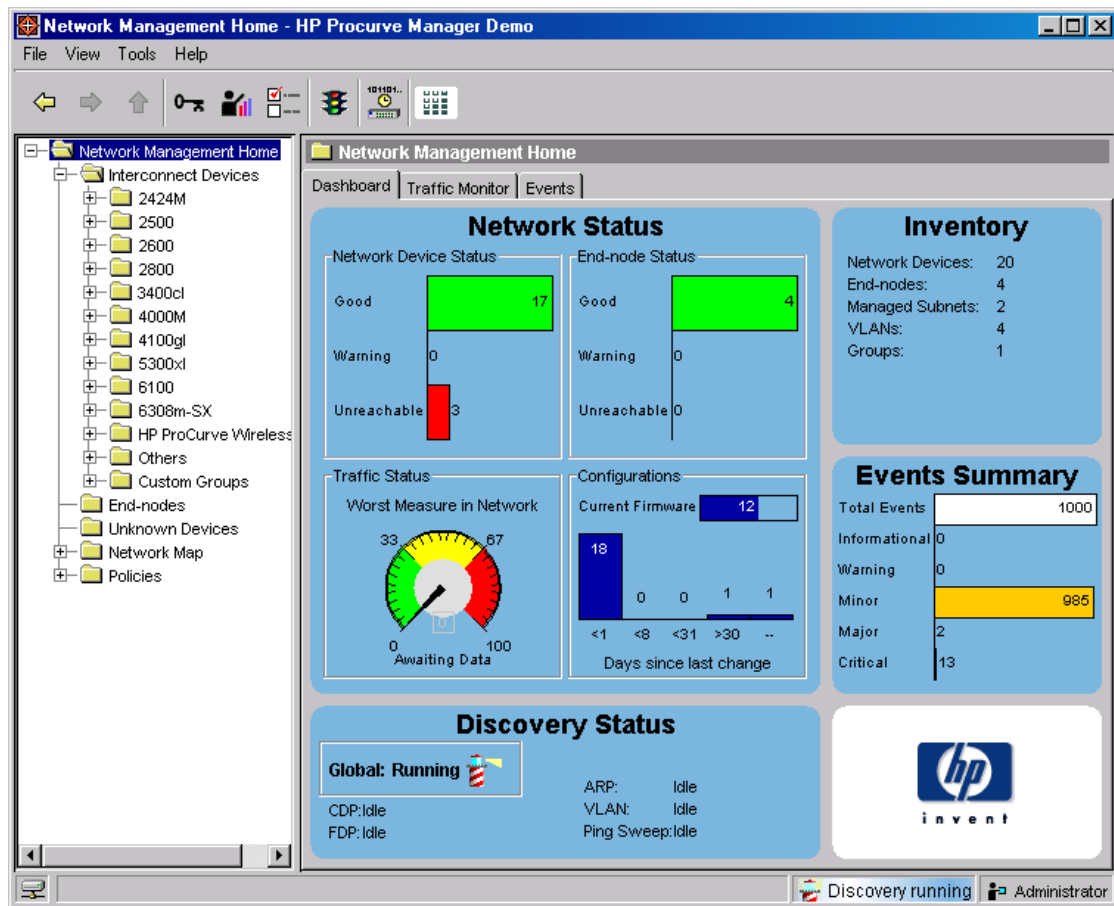


Figure 2-4. Example of the Home Page for ProCurve Manager Plus

PCM and PCM+ enable greater control, uptime, and performance in your network:

- Features and benefits of ProCurve Manager:
 - **Network Status Summary:** Upon boot-up, a network status screen displays high-level information on network devices, end nodes, events, and traffic levels. From here, users can research any one of these areas to get more details.
 - **Alerts and Troubleshooting:** An events summary screen displays alerts to the user and categorizes them by severity, making it easier to track where bottlenecks and issues exist in the network. Alerts present detailed information on the problem, even down to the specific port.
 - **Automatic Device Discovery:** This feature is customized for fast discovery of all ProCurve manageable network devices. The user can define which IP subnets to discover.
 - **Topology and Mapping:** This feature automatically creates a map of discovered network devices. Maps are color-coded to reflect device status and can be viewed at multiple levels (physical view, subnet view, or VLAN view).
 - **Device Management:** Many device-focused tasks can be performed directly by the software, or the user can access web-browser and command-line interfaces with the click of a button to manage individual devices from inside the tool.
- Features and benefits of ProCurve Manager Plus:
 - **All of the Features of ProCurve Manager:** Refer to the above listing.
 - **In-Depth Traffic Analysis:** An integrated, low-overhead traffic monitor interface shows detailed information on traffic throughout the network. Using enhanced traffic analysis protocols such as Extended RMON and sFlow, users can monitor overall traffic levels, segments with the highest traffic, or even the top users within a network segment.
 - **Group and Policy Management:** Changes in configuration are tracked and logged, and archived configurations can be applied to one or many devices. Configurations can be compared over time or between two devices, with the differences highlighted for users.
 - **Advanced VLAN Management:** A new, easy-to-use VLAN management interface allows users to create and assign VLANs across the entire network, without having to access each network device individually.

- **Device Software Updates:** This feature automatically obtains new device software images from ProCurve and updates devices, allowing users to download the latest version or choose the desired version. Updates can be scheduled easily across large groups of devices, all at user-specified times.
- **Investment Protection:** The modular software architecture of ProCurve Manager Plus will allow ProCurve to offer network administrators add-on software solutions that complement their needs.

Selecting a Management Interface

Advantages of Using ProCurve Manager or ProCurve Manager Plus

— This page is intentionally unused. —

Using the Menu Interface

Contents

- Overview** 3-2
- Starting and Ending a Menu Session** 3-3
 - How To Start a Menu Interface Session 3-4
 - How To End a Menu Session and Exit from the Console: 3-5
- Main Menu Features** 3-7
- Screen Structure and Navigation** 3-9
- Rebooting the Switch** 3-12
- Menu Features List** 3-14
- Where To Go From Here** 3-15

Overview

This chapter describes the following features:

- Overview of the Menu Interface (page 3-2)
- Starting and ending a Menu session (page 3-3)
- The Main Menu (page 3-7)
- Screen structure and navigation (page 3-9)
- Rebooting the switch (page 3-12)

The menu interface operates through the switch console to provide you with a subset of switch commands in an easy-to-use menu format enabling you to:

- Perform a “quick configuration” of basic parameters, such as the IP addressing needed to provide management access through your network
- Configure these features:
 - Manager and Operator passwords
 - System parameters
 - IP addressing
 - Time protocol
 - Ports
 - Trunk groups
 - A network monitoring port
 - Stack Management (3400cl and 6400cl switches only)
 - Spanning Tree operation
 - SNMP community names
 - IP authorized managers
 - VLANs (Virtual LANs) and GVRP
- View status, counters, and Event Log information
- Update switch software
- Reboot the switch

For a detailed list of menu features, see the “Menu Features List” on page 3-14.

Privilege Levels and Password Security. ProCurve strongly recommends that you configure a Manager password to help prevent unauthorized access to your network. A Manager password grants full read-write access to the switch. An Operator password, if configured, grants access to status and counter, Event Log, and the Operator level in the CLI. After you configure passwords on the switch and log off of the interface, access to the menu interface (and the CLI and web browser interface) will require entry of either the Manager or Operator password. (If the switch has only a Manager password, then someone without a password can still gain read-only access.)

Note

If the switch has neither a Manager nor an Operator password, anyone having access to the console interface can operate the console with full manager privileges. Also, if you configure only an Operator password, entering the Operator password enables full manager privileges.

For more information on passwords, refer to the *Access Security Guide* for your switch.

Menu Interaction with Other Interfaces.

- The menu interface displays the current running-config parameter settings. You can use the menu interface to save configuration changes made in the CLI only if the CLI changes are in the running config when you save changes made in the menu interface. (For more on how switch memory manages configuration changes, see Chapter 6, “Switch Memory and Configuration”.)
- A configuration change made through any switch interface overwrites earlier changes made through any other interface.
- The Menu Interface and the CLI (Command Line Interface) both use the switch console. To enter the menu from the CLI, use the **menu** command. To enter the CLI from the Menu interface, select **Command Line (CLI)** option.)

Starting and Ending a Menu Session

You can access the menu interface using any of the following:

- A direct serial connection to the switch’s console port, as described in the installation guide you received with the switch
- A Telnet connection to the switch console from a networked PC or the switch’s web browser interface. Telnet requires that an IP address and subnet mask compatible with your network have already been configured on the switch.

Note

This section assumes that either a terminal device is already configured and connected to the switch (see the *Installation and Getting Started Guide* shipped with your switch) or that you have already configured an IP address on the switch (required for Telnet access).

How To Start a Menu Interface Session

In its factory default configuration, the switch console starts with the CLI prompt. To use the menu interface with Manager privileges, go to the Manager level prompt and enter the **menu** command.

1. Use one of these methods to connect to the switch:
 - A PC terminal emulator or terminal
 - Telnet
2. Do one of the following:
 - If you are using Telnet, go to step 3.
 - If you are using a PC terminal emulator or a terminal, press **[Enter]** one or more times until a prompt appears.
3. When the switch screen appears, do one of the following:
 - If a password has been configured, the password prompt appears.

```
Password: _
```

Type the Manager password and press **[Enter]**. Entering the Manager password gives you manager-level access to the switch. (Entering the Operator password gives you operator-level access to the switch. Refer to the *Access Security Guide* for your switch.)
 - If no password has been configured, the CLI prompt appears. Go to the next step.
4. When the CLI prompt appears, display the Menu interface by entering the **menu** command. For example:

```
ProCurve# menu [Enter]
```

results in the following display:

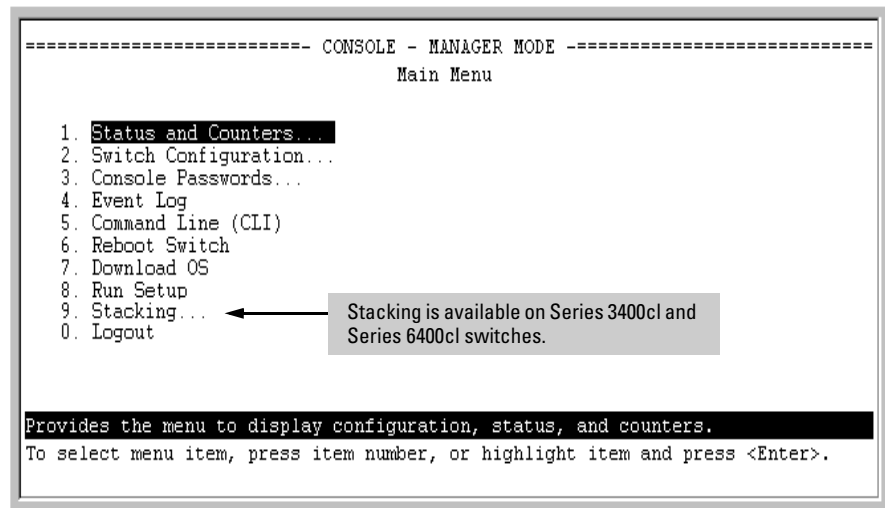


Figure 3-1. Example of the Main Menu with Manager Privileges

For a description of Main Menu features, see “Main Menu Features” on page 3-7.

Note

To configure the switch to start with the menu interface instead of the CLI, go to the Manager level prompt in the CLI, enter the **setup** command, and in the resulting display, change the **Logon Default** parameter to **Menu**. For more information, see the *Installation and Getting Started Guide* you received with the switch.

How To End a Menu Session and Exit from the Console

The method for ending a menu session and exiting from the console depends on whether, during the session, you made any changes to the switch configuration that require a switch reboot to activate. (Most changes via the menu interface need only a **Save**, and do not require a switch reboot.) Configuration changes needing a reboot are marked with an asterisk (*) next to the configured item in the menu and also next to the **Switch Configuration** item in the Main Menu.

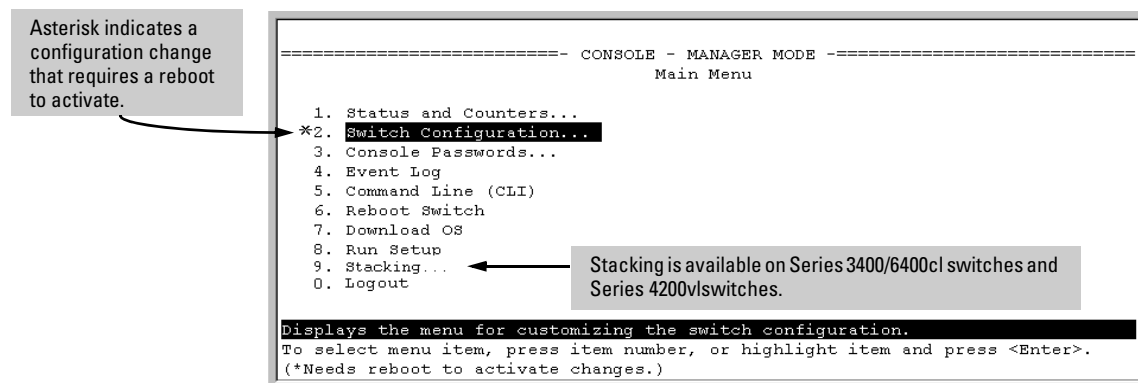


Figure 3-2. Example Indication of a Configuration Change Requiring a Reboot

1. In the current session, if you have not made configuration changes that require a switch reboot to activate, return to the Main Menu and press **[0]** (zero) to log out. Then just exit from the terminal program, turn off the terminal, or quit the Telnet session.
2. If you *have* made configuration changes that require a switch reboot—that is, if an asterisk (*) appears next to a configured item or next to **Switch Configuration** in the Main Menu:
 - a. Return to the Main Menu.
 - b. Press **[6]** to select **Reboot Switch** and follow the instructions on the reboot screen.

Rebooting the switch terminates the menu session, and, if you are using Telnet, disconnects the Telnet session.

(See “Rebooting To Activate Configuration Changes” on page 3-13.)

3. Exit from the terminal program, turn off the terminal, or close the Telnet application program.

Main Menu Features

```
===== CONSOLE - MANAGER MODE =====
                          Main Menu

1. Status and Counters...
2. Switch Configuration...
3. Console Passwords...
4. Event Log
5. Command Line (CLI)
6. Reboot Switch
7. Download OS
8. Run Setup
9. Stacking...
0. Logout

Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 3-3. The Main Menu View with Manager Privileges

The Main Menu gives you access to these Menu interface features:

- **Status and Counters:** Provides access to display screens showing switch information, port status and counters, port and VLAN address tables, and spanning tree information. (See Appendix B, “Monitoring and Analyzing Switch Operation”.)
- **Switch Configuration:** Provides access to configuration screens for displaying and changing the current configuration settings. (See the Contents listing at the front of this manual.) For a listing of features and parameters configurable through the menu interface, see the “Menu Features List” on page 3-14 .
- **Console Passwords:** Provides access to the screen used to set or change Manager-level and Operator-level passwords, and to delete Manager and Operator password protection. (Refer to the chapter on configuring usernames and passwords in the *Access Security Guide* for your switch.)
- **Event Log:** Enables you to read progress and error messages that are useful for checking and troubleshooting switch operation. (See “Using the Event Log To Identify Problem Sources” on page C-27.)

- **Command Line (CLI):** Selects the Command Line Interface at the same level (Manager or Operator) that you are accessing in the Menu interface. (Refer to chapter 3, “Using the Command Line Interface (CLI)”.)
- **Reboot Switch:** Performs a “warm” reboot of the switch, which clears most temporary error conditions, resets the network activity counters to zero, and resets the system up-time to zero. A reboot is required to activate a change in the VLAN Support parameter. (See “Rebooting from the Menu Interface” on page 6-11.)
- **Download OS:** Enables you to download a new switch software version to the switch. (See Appendix A, “File Transfers”.)
- **Run Setup:** Displays the Switch Setup screen for quickly configuring basic switch parameters such as IP addressing, default gateway, logon default interface, spanning tree, and others. (See the *Installation and Getting Started Guide* for your switch.)
- **Logout:** Closes the Menu interface and console session, and disconnects Telnet access to the switch. (See “How to End a Menu Session and Exit from the Console” on page 3-5.)

Screen Structure and Navigation

Menu interface screens include these three elements:

- Parameter fields and/or read-only information such as statistics
- Navigation and configuration actions, such as Save, Edit, and Cancel
- Help line to describe navigation options, individual parameters, and read-only data

For example, in the following System Information screen:

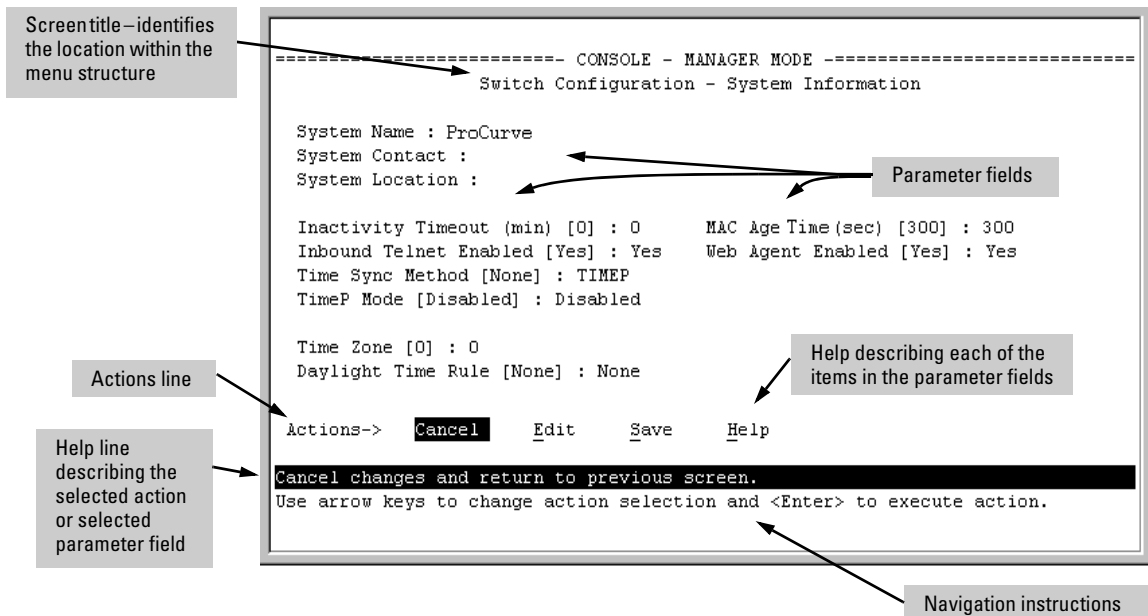


Figure 3-4. Elements of the Screen Structure

“Forms” Design. The configuration screens, in particular, operate similarly to a number of PC applications that use forms for data entry. When you first enter these screens, you see the current configuration for the item you have selected. To change the configuration, the basic operation is to:

1. Press **[E]** to select the **Edit** action.
2. Navigate through the screen making all the necessary configuration changes. (See Table 4-1 on the next page.)
3. Press **[Enter]** to return to the **Actions** line. From there you can save the configuration changes or cancel the changes. Cancel returns the configuration to the values you saw when you first entered the screen.

Table 3-1. How To Navigate in the Menu Interface

Task:	Actions:
Execute an action from the “Actions →” list at the bottom of the screen:	<p>Use either of the following methods:</p> <ul style="list-style-type: none"> • Use the arrow keys (← , or →) to highlight the action you want to execute, then press [Enter]. • Press the key corresponding to the capital letter in the action name. For example, in a configuration menu, press [E] to select Edit and begin editing parameter values.
Reconfigure (edit) a parameter setting or a field:	<ol style="list-style-type: none"> 1. Select a configuration item, such as System Name. (See figure 3-4.) 2. Press [E] (for Edit on the Actions line). 3. Use [Tab] or the arrow keys (← , → , ↑ , or ↓) to highlight the item or field. 4. Do one of the following: <ul style="list-style-type: none"> – If the parameter has preconfigured values, either use the Space bar to select a new option or type the first part of your selection and the rest of the selection appears automatically. (The help line instructs you to “Select” a value.) – If there are no preconfigured values, type in a value (the Help line instructs you to “Enter” a value). 5. If you want to change another parameter value, return to step 3. 6. If you are finished editing parameters in the displayed screen, press [Enter] to return to the Actions line and do one of the following: <ul style="list-style-type: none"> – To save and activate configuration changes, press [S] (for the Save action). This saves the changes in the startup configuration and also implements the change in the currently running configuration. (See Chapter 6, “Switch Memory and Configuration”.) – To exit from the screen without saving any changes that you have made (or if you have not made changes), press [C] (for the Cancel action). <p>Note: In the menu interface, executing Save activates most parameter changes and saves them in the startup configuration (or flash) memory, and it is therefore not necessary to reboot the switch after making these changes. But if an asterisk appears next to any menu item you reconfigure, the switch will not activate or save the change for that item until you reboot the switch. In this case, rebooting should be done after you have made all desired changes and then returned to the Main Menu.</p> <ol style="list-style-type: none"> 7. When you finish editing parameters, return to the Main Menu. 8. If necessary, reboot the switch by highlighting Reboot Switch in the Main Menu and pressing [Enter]. (See the Note, above.)
Exit from a read-only screen.	Press [B] (for the Back action).

To get Help on individual parameter descriptions. In most screens there is a **Help** option in the **Actions** line. Whenever any of the items in the **Actions** line is highlighted, press **[H]**, and a separate help screen is displayed. For example:

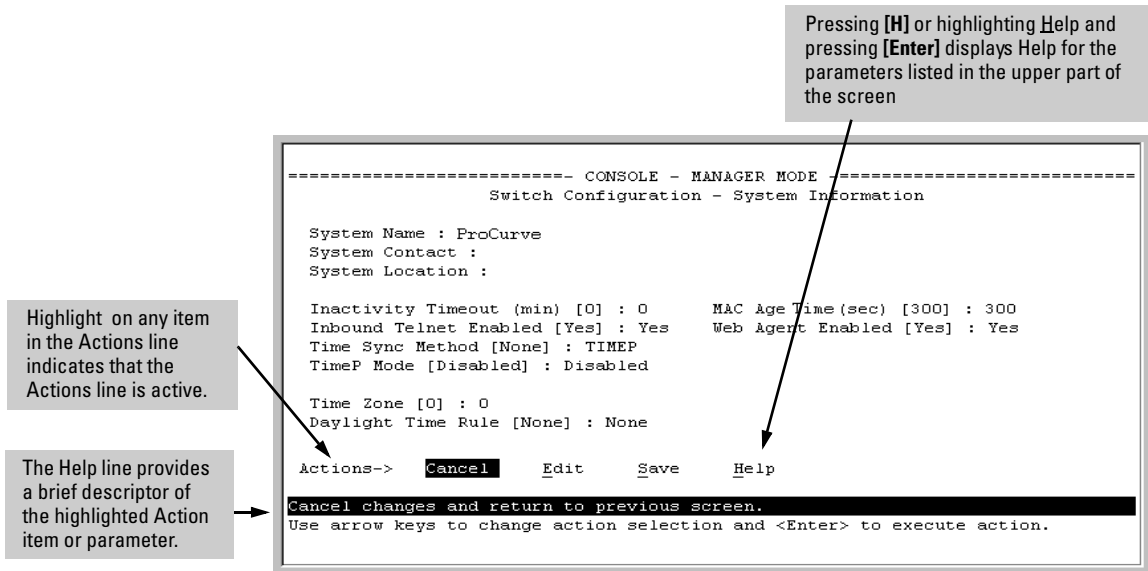


Figure 3-5. Example Showing How To Display Help

To get Help on the actions or data fields in each screen: Use the arrow keys (**←**, **→**, **↑**, or **↓**) to select an action or data field. The help line under the **Actions** items describes the currently selected action or data field.

For guidance on how to navigate in a screen: See the instructions provided at the bottom of the screen, or refer to "Screen Structure and Navigation" on page 3-9.)

Rebooting the Switch

Rebooting the switch from the menu interface

- Terminates all current sessions and performs a reset of the operating system
- Activates any menu interface configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that **Reboot Switch** is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)

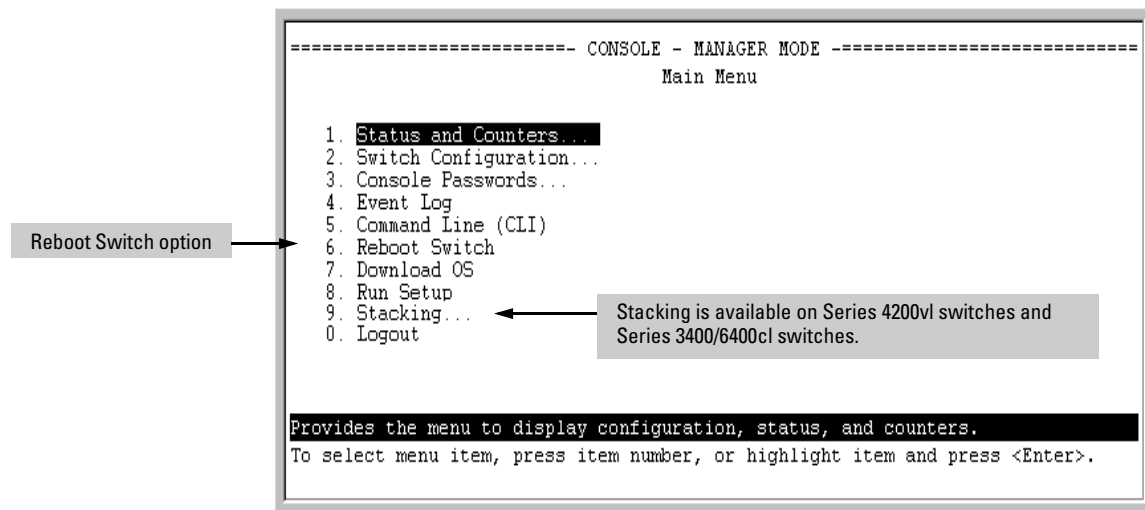


Figure 3-6. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters in the menu interface become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support parameter**. (To access this parameter, go to the Main Menu and select:

2. Switch Configuration

8. VLAN Menu

1. VLAN Support.

If you make configuration changes in the menu interface that require a reboot, the switch displays an asterisk (*) next to the menu item in which the change has been made. For example, if you change and save the value for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** ..entry in the Main Menu, as shown in figure 4-6:

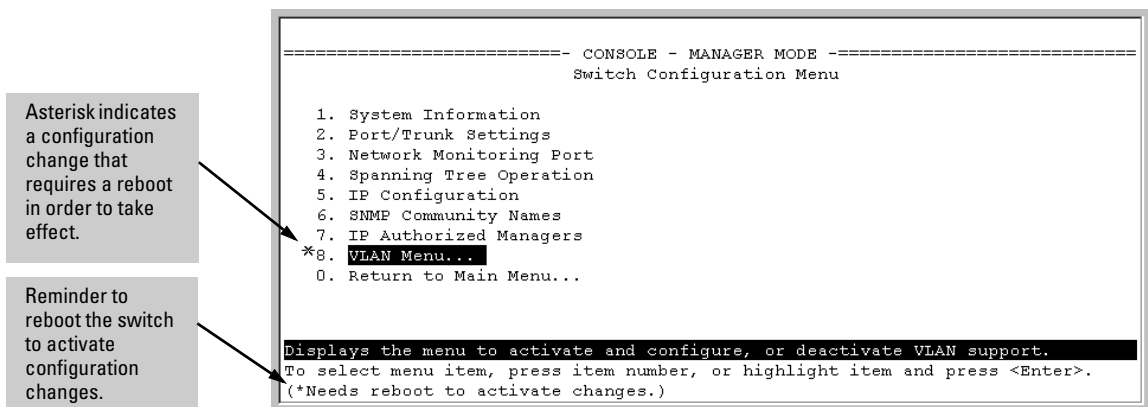


Figure 3-7. Indication of a Configuration Change Requiring a Reboot

To activate changes indicated by the asterisk, go to the Main Menu and select the **Reboot Switch** option.

Note

Executing the **write memory** command in the CLI does not affect pending configuration changes indicated by an asterisk in the menu interface. That is, only a reboot from the menu interface or a **boot** or **reload** command from the CLI will activate a pending configuration change indicated by an asterisk.

Menu Features List

Status and Counters

- General System Information
- Switch Management Address Information
- Port Status
- Port Counters
- Address Table
- Port Address Table
- Spanning Tree Information

Switch Configuration

- System Information
- Port/Trunk Settings
- Network Monitoring Port
- Spanning Tree Operation
- IP Configuration
- SNMP Community Names
- IP authorized Managers
- VLAN Menu

Console Passwords

Event Log

Command Line (CLI)

Reboot Switch

Download OS (Download Switch Software)

Run Setup

Stacking (Series 3400cl, Series 6400cl and Series 4200vl switches)

- Stacking Status (This Switch)
- Stacking Status (All)
- Stack Configuration
- Stack Management (*Stack Commander Only*)
- Stack Access (*Stack Commander Only*)

Logout

Where To Go From Here

This chapter provides an overview of the menu interface and how to use it. The following table indicates where to turn for detailed information on how to use the individual features available through the menu interface.

Option:	Turn to:
To use the Run Setup option	Refer to the <i>Installation and Getting Started Guide</i> shipped with the switch.
To view and monitor switch status and counters	Appendix B, "Monitoring and Analyzing Switch Operation"
To learn how to configure and use passwords and other security features	Refer to the <i>Access Security Guide</i> for your switch.
To learn how to use the Event Log	"Using the Event Log To Identify Problem Sources" on page C-27
To learn how the CLI operates	Chapter 4, "Using the Command Line Interface (CLI)"
To download switch software	Appendix A, "File Transfers"
For a description of how switch memory handles configuration changes	Chapter 6, "Switch Memory and Configuration"
For information on other switch features and how to configure them	Refer to the Table of Contents at the front of this guide, and to "Sources for More Information" on page 1-4.

— This page is intentionally unused. —

Using the Command Line Interface (CLI)

Contents

Overview	4-2
Accessing the CLI	4-2
Using the CLI	4-2
Privilege Levels at Logon	4-3
Privilege Level Operation	4-4
Operator Privileges	4-4
Manager Privileges	4-5
How To Move Between Levels	4-7
Listing Commands and Command Options	4-8
Listing Commands Available at Any Privilege Level	4-8
Listing Command Options	4-10
Displaying CLI “Help”	4-11
Configuration Commands and the Context Configuration Modes ..	4-12
Configuring Custom Login Banners for the Console and Web Browser Interfaces	4-15
Banner Operation with Telnet, Serial, or SSHv2 Access	4-16
Banner Operation with Web Browser Access	4-16
Configuring and Displaying a Non-Default Banner	4-16
Example of Configuring and Displaying a Banner	4-17
Operating Notes	4-19
CLI Control and Editing	4-21

Overview

The CLI is a text-based command interface for configuring and monitoring the switch. The CLI gives you access to the switch's full set of commands while providing the same password protection that is used in the web browser interface and the menu interface.

Accessing the CLI

Like the menu interface, the CLI is accessed through the switch console, and in the switch's factory default state, is the default interface when you start a console session. You can access the console out-of-band by directly connecting a terminal device to the switch, or in-band by using Telnet either from a terminal device or through the web browser interface.

Also, if you are using the menu interface, you can access the CLI by selecting the **Command Line (CLI)** option in the Main Menu.

Using the CLI

The CLI offers these privilege levels to help protect the switch from unauthorized access:

1. Operator
2. Manager
3. Global Configuration
4. Context Configuration

Note

CLI commands are not case-sensitive.

When you use the CLI to make a configuration change, the switch writes the change to the Running-Config file in volatile memory. This allows you to test your configuration changes before making them permanent. To make changes permanent, you must use the **write memory** command to save them to the Startup-Config file in non-volatile memory. If you reboot the switch without first using **write memory**, all changes made since the last reboot or **write memory** (whichever is later) will be lost. For more on switch memory and saving configuration changes, see Chapter 6, “Switch Memory and Configuration”.

Privilege Levels at Logon

Privilege levels control the type of access to the CLI. To implement this control, you must set at least a Manager password. *Without a Manager password configured, anyone having serial port, Telnet, or web browser access to the switch can reach all CLI levels.* (For more on setting passwords, refer to the chapter on usernames and passwords in the *Access Security Guide* for your switch.)

When you use the CLI to log on to the switch, and passwords are set, you will be prompted to enter a password. For example:

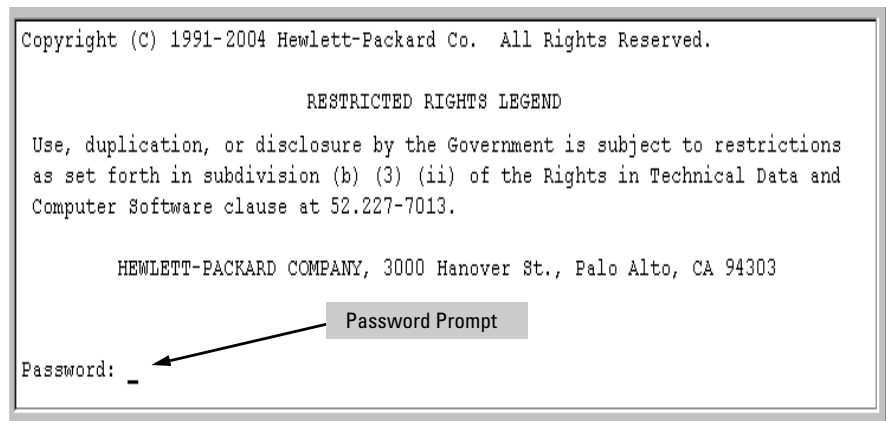


Figure 4-1. Example of CLI Log-On Screen with Password(s) Set

In the above case, you will enter the CLI at the level corresponding to the password you provide (operator or manager).

If no passwords are set when you log onto the CLI, you will enter at the Manager level. For example:

ProCurve# _

Caution

ProCurve strongly recommends that you configure a Manager password. If a Manager password is not configured, then the Manager level is not password-protected, and anyone having in-band or out-of-band access to the switch may be able to reach the Manager level and compromise switch and network security. Note that configuring only an Operator password *does not* prevent access to the Manager level by intruders who have the Operator password.

Pressing the Clear button on the front of the switch removes password protection. *For this reason, it is recommended that you protect the switch from physical access by unauthorized persons.* If you are concerned about switch security and operation, you should install the switch in a secure location, such as a locked wiring closet.

Privilege Level Operation

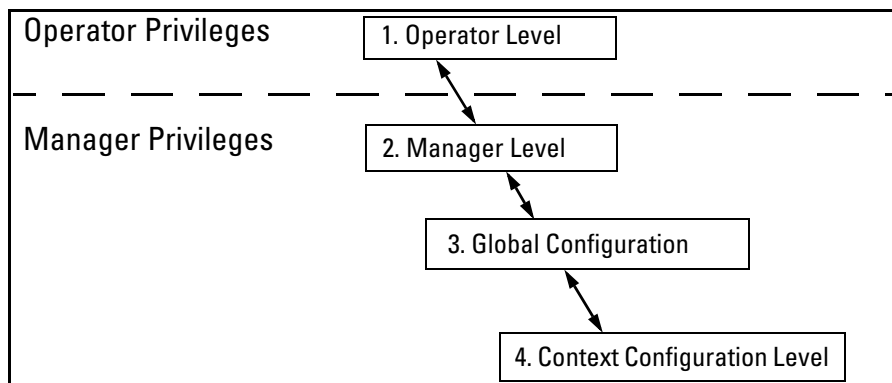


Figure 4-2. Access Sequence for Privilege Levels

Operator Privileges

At the Operator level you can examine the current configuration and move between interfaces without being able to change the configuration. A ">" character delimits the Operator-level prompt. For example:

ProCurve> _ (Example of the Operator prompt.)

When using **enable** to move to the Manager level, the switch prompts you for the Manager password if one has already been configured.

Manager Privileges

Manager privileges give you three additional levels of access: Manager, Global Configuration, and Context Configuration. (See figure.) A “#” character delimits any Manager prompt. For example:

ProCurve#_ *Example of the Manager prompt.*

- **Manager level:** Provides all Operator level privileges plus the ability to perform system-level actions that do not require saving changes to the system configuration file. The prompt for the Manager level contains only the system name and the “#” delimiter, as shown above. To select this level, enter the **enable** command at the Operator prompt and enter the Manager password, when prompted. For example:

```
ProCurve> enable      Enter enable at the Operator prompt.
Password:             CLI prompt for the Manager password.
ProCurve# _           The Manager prompt appears after the
                       correct Manager password is entered.
```

- **Global Configuration level:** Provides all Operator and Manager level privileges, and enables you to make configuration changes to any of the switch’s software features. The prompt for the Global Configuration level includes the system name and “(config)”. To select this level, enter the **config** command at the Manager prompt. For example:

```
HPswitch# config      Enter config at the Manager prompt.
HPswitch(config)#_    The Global Config prompt.
```

- **Context Configuration level:** Provides all Operator and Manager privileges, and enables you to make configuration changes in a specific context, such as one or more ports or a VLAN. The prompt for the Context Configuration level includes the system name and the selected context. For example:

```
ProCurve(eth-1)#
ProCurve(vlan-10)#
```

The Context level is useful, for example, for executing several commands directed at the same port or VLAN, or if you want to shorten the command strings for a specific context area. To select this level, enter the specific context at the Global Configuration level prompt. For example, to select the context level for an existing VLAN with the VLAN ID of 10, you would enter the following command and see the indicated result:

```
ProCurve(config)# vlan 10
ProCurve(vlan-10)#
```

Table 4-1. Privilege Level Hierarchy

Privilege Level	Example of Prompt and Permitted Operations		
Operator Privilege			
Operator Level	ProCurve>	show < <i>command</i> > setup	View status and configuration information.
		ping < <i>argument</i> > link-test < <i>argument</i> >	Perform connectivity tests.
		enable	Move from the Operator level to the Manager level.
		menu	Move from the CLI interface to the menu interface.
		logout	Exit from the CLI interface and terminate the console session.
		exit	Terminate the current session (same as logout).
Manager Privilege			
Manager Level	ProCurve#		Perform system-level actions such as system control, monitoring, and diagnostic commands, plus any of the Operator-level commands. For a list of available commands, enter ? at the prompt.
Global Configuration Level	ProCurve (config) #		Execute configuration commands, plus all Operator and Manager commands. For a list of available commands, enter ? at the prompt.
Context Configuration Level	ProCurve (eth-5) # ProCurve (vlan-100) #		Execute context-specific configuration commands, such as a particular VLAN or switch port. This is useful for shortening the command strings you type, and for entering a series of commands for the same context. For a list of available commands, enter ? at the prompt.

How To Move Between Levels

Change in Levels	Example of Prompt, Command, and Result	
Operator level to Manager level	ProCurve> enable Password: _ ProCurve# _	After you enter enable , the Password prompt appears. After you enter the Manager password, the system prompt appears with the # symbol:
Manager level to Global configuration level	ProCurve# config ProCurve(config)#	
Global configuration level to a Context configuration level	ProCurve(config)# vlan 10 ProCurve(vlan-10)#	
Context configuration level to another Context configuration level	ProCurve(vlan-10)# interface e 3 ProCurve(int-3)#	The CLI accepts "e" as the abbreviated form of "ethernet".
Move from any level to the preceding level	ProCurve(int-3)# exit ProCurve(config)# exit ProCurve# exit ProCurve>	
Move from any level to the Manager level	ProCurve(int-3)# end ProCurve# -or- ProCurve(config)# end ProCurve#	

Moving Between the CLI and the Menu Interface. When moving between interfaces, the switch retains the current privilege level (Manager or Operator). That is, if you are at the Operator level in the menu and select the **Command Line Interface (CLI)** option from the Main Menu, the CLI prompt appears at the Operator level.

Changing Parameter Settings. Regardless of which interface is used (CLI, menu interface, or web browser interface), the most recently configured version of a parameter setting overrides any earlier settings for that parameter.

For example, if you use the menu interface to configure an IP address of “X” for VLAN 1 and later use the CLI to configure a different IP address of “Y” for VLAN 1, then “Y” replaces “X” as the IP address for VLAN 1 in the running-config file. If you subsequently execute **write memory** in the CLI, then the switch also stores “Y” as the IP address for VLAN 1 in the startup-config file. (For more on the startup-config and running config files, see Chapter 6, “Switch Memory and Configuration”).

Listing Commands and Command Options

At any privilege level you can:

- List all of the commands available at that level
- List the options for a specific command

Listing Commands Available at Any Privilege Level

At a given privilege level you can list and execute the commands that level offers, plus all of the commands available at preceding levels. For example, at the Operator level, you can list and execute only the Operator level commands. However, at the Manager level, you can list and execute the commands available at both the Operator and Manager levels.

Type “?” To List Available Commands. 1. Typing the ? symbol lists the commands you can execute at the current privilege level. For example, typing ? at the Operator level produces this listing:

```
ProCurve> ?
enable
exit
link-test
logout
menu
ping
show
traceroute
HPswitch>
```

Figure 4-3. Example of the Operator Level Command Listing

Typing ? at the Manager level produces this listing:

```
ProCurve# ?
boot          Reboot the device.
clear         Clear table/statistics or authorized client public
              keys.
configure     Enter the Configuration context.
copy          Copy datafiles to/from the switch.
debug         Enable/disable debug logging.
display       Display the running/saved configuration.
end           Return to the Manager Exec context.
erase         Erase the configuration file stored in flash or.
getMIB        Retrieve and display the value of the MIB objects
              specified.
kill          Kill other active console, telnet, or ssh sessions.
log           Display log events.
page          Toggle paging mode.
print         Execute a command and redirect its output to the device
              channel for current session.
redo          Re-execute a command from history.
reload        Warm reboot of the switch.
repeat        Repeat execution of a previous command.
setMIB        Set the value of a MIB object.
setup         Enter the 'Switch Setup' screen for basic switch
              configuration.
-- MORE --    --, next page: Space, next line: Enter, quit: Control-C
```

When -- MORE -- appears, use the Space bar or [Return] to list additional commands.

Figure 4-4.Example of the Manager-Level Command Listing

When -- **MORE** -- appears, there are more commands in the listing. To list the next screenfull of commands, press the Space bar. To list the remaining commands one-by-one, repeatedly press [Enter].

Typing ? at the Global Configuration level or the Context Configuration level produces similar results.

Use [Tab] To Search for or Complete a Command Word. You can use [Tab] to help you find CLI commands or to quickly complete the current word in a command. To do so, type one or more consecutive characters in a command and then press [Tab] (with no spaces allowed). For example, at the Global Configuration level, if you press [Tab] immediately after typing “t”, the CLI displays the available command options that begin with “t”. For example:

```
ProCurve(config)# t [Tab]
tacacs-server
telnet-server
time
timesync
trunk
telnet
terminal
traceroute
ProCurve(config)# t
```

As mentioned above, if you type part of a command word and press **[Tab]**, the CLI completes the current word (if you have typed enough of the word for the CLI to distinguish it from other possibilities), including hyphenated extensions. For example:

```
ProCurve(config)# port-[Tab]
ProCurve(config)# port-security _
```

Pressing **[Tab]** after a completed command word lists the further options for that command.

```
ProCurve(config)# qos [Tab]

udp-portSet UDP port based priority.
tcp-portSet TCP port based priority.
device-priorityConfigure device-based priority.
dscp-mapDefine mapping between a DSCP
(Differentiated-Services Codepoint)
value and 802.1p priority.
type-of-serviceConfigure the Type-of-Service
method the device uses to
prioritize IP traffic.
```

Listing Command Options

You can use the CLI to remind you of the options available for a command by entering command keywords followed by **?**. For example, suppose you want to see the command options for configuring the console settings:

This example displays the command options for configuring the switch's console settings.

```
(ProCurve(config)# _console_)
  baud-rate      Set the data transmission speed for the device connect
                  sessions initiated through the Console port.
  events         Set level of the events displayed in the device's Events
                  Log.
  flow-control    Set the Flow Control Method: default is xon-xoff.
  inactivity-timer Set the number of minutes of no activity detected on the
                  Console port before the switch terminates a
                  communication session.
  screen-refresh  Set default number of seconds before screen is refreshed
                  on the repeat command.
  terminal        Set type of terminal being used (default is vt100).
```

Figure 4-5. Example of How To List the Options for a Specific Command

Displaying CLI “Help”

CLI Help provides two types of context-sensitive information:

- Command list with a brief summary of each command’s purpose
- Detailed information on how to use individual commands

Displaying Command-List Help.

Syntax: help

*Displays a listing of command Help summaries for all commands available at the current privilege level. That is, at the Operator level, executing **help** displays the Help summaries only for Operator-Level commands. At the Manager level, executing **help** displays the Help summaries for both the Operator and Manager levels, and so on.*

For example, to list the Operator-Level commands with their purposes:

```
ProCurve> help

enable          Enter the Manager Exec context.
exit            Return to the previous context or terminate current
               console/telnet session if you are in the Operator
               context level.
link-test       Test the connection to a MAC address on the LAN.
logout          Terminate this console/telnet session.
menu            Change console user interface to menu system.
ping            Send IP Ping requests to a device on the network.
show            Display switch operation information.
traceroute      Send traceroute to a device on the network.
```

Figure 4-6. Example of Context-Sensitive Command-List Help

Displaying Help for an Individual Command.

Syntax: < command-string > help

This option displays Help for any command available at the current context level.

For example, to list the Help for the **interface** command in the Global Configuration privilege level:

```
ProCurve(config)# interface help
Usage: [no] interface [ethernet] PORT-LIST [...]

Description: Enter the Interface Configuration Level, or execute one
             command for that level. Without optional parameters
             specified, the 'interface' command changes the context to
             the Interface Configuration Context Level for execution of
             configuration changes to the port or ports in the PORT-LIST.
             The 'interface [ethernet] PORT-LIST' can be followed by any
             command from the Interface Configuration Context Level in the
             same command line. In this case the context level is not
             changed, but the command is also executed for the port or ports
             in the PORT-LIST. Use 'interface [ethernet] PORT-LIST ?'
             to get a list of all valid commands.
```

Figure 4-7.Example of How To Display Help for a Specific Command

Note that trying to list the help for an individual command from a privilege level that does not include that command results in an error message. For example, trying to list the help for the **interface** command while at the global configuration level produces this result:

```
ProCurve# speed-duplex help
Invalid input: speed-duplex
```

Configuration Commands and the Context Configuration Modes

You can execute any configuration command in the global configuration mode or in selected context modes. However, using a context mode enables you to execute context-specific commands faster, with shorter command strings.

The switch offers interface (port or trunk group) and VLAN context configuration modes:

Port or Trunk-Group Context . Includes port- or trunk-specific commands that apply only to the selected port(s) or trunk group, plus the global configuration, Manager, and Operator commands. The prompt for this mode includes the identity of the selected port(s):

```
ProCurve(config)# interface c3-c6
ProCurve(eth-C5-C8)#

ProCurve(config)# interface trk1
ProCurve(eth-Trk1)#
```

*Commands executed at configuration level for entering port and **trk1** static trunk-group contexts, and resulting prompts showing port or static trunk contexts..*

```
ProCurve(eth-C5-C8) #
ProCurve(eth-Trk1) #

ProCurve(eth-C5-C8) # ?
ProCurve(eth-C5-C8) # ?
```

Lists the commands you can use in the port or static trunk context, plus the Manager, Operator, and context commands you can execute at this level.

In the port context, the first block of commands in the "?" listing show the context-specific commands that will affect only ports C3-C6.

```

ProCurve(eth-3-6) # ?
| broadcast-limit  Set a broadcast traffic percentage limit.
| disable          Disable port(s).
| enable          Enable port(s).
| flow-control     Enable/disable flow control on the port(s).
| gvrp            Set the GVRP timers on the port (hundreths of a
|                second).
| lacp            Define whether LACP is enabled on the port, and whether
|                it is in active or passive mode when enabled.
| mdix-mode        Set port MDI/MDIX mode (default: auto).
| monitor         Define either the port is to be monitored or not.
| name            Set/unset a name for the port(s).
| qos             Set port-based priority.
| rate-limit      Enable/disable and configure rate-limiting for incoming
|                traffic on the port(s).
| speed-duplex    Define mode of operation for the port(s).
| unknown-vlans   Configure GVRP on the port(s).
-- MORE --, next page: Space, next line: Enter, quit: Control-C

```

The remaining commands in the listing are Manager, Operator, and context commands.

Figure 4-8. Context-Specific Commands Affecting Port Context

VLAN Context . Includes VLAN-specific commands that apply only to the selected VLAN, plus Manager and Operator commands. The prompt for this mode includes the VLAN ID of the selected VLAN. For example, if you had already configured a VLAN with an ID of 100 in the switch:

```
ProCurve(config)# vlan 100
```

Command executed at configuration level to enter VLAN 100 context.

```
ProCurve(vlan-100)#
```

Resulting prompt showing VLAN 100 context.

```
ProCurve(vlan-100)# ?
```

Lists commands you can use in the VLAN context, plus Manager, Operator, and context commands you can execute at this level.

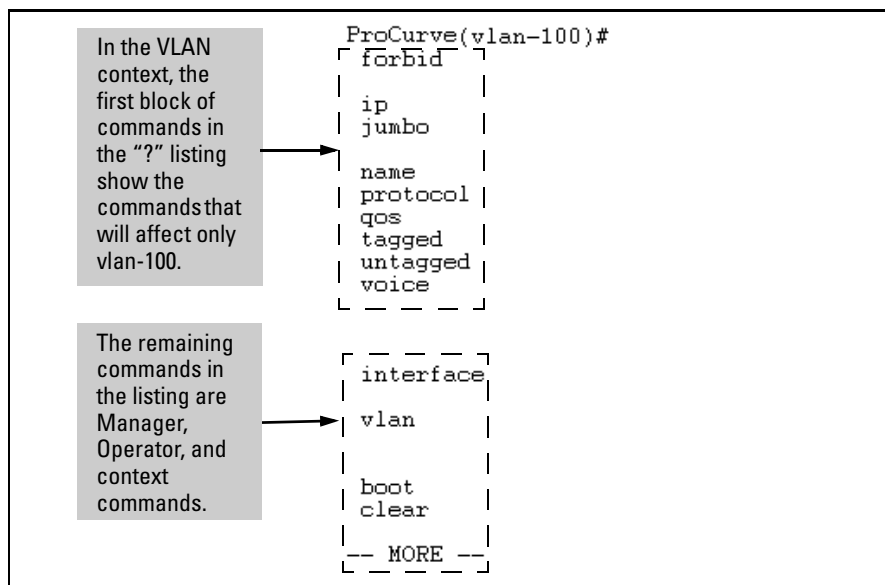


Figure 4-9. Context-Specific Commands Affecting VLAN Context

Configuring Custom Login Banners for the Console and Web Browser Interfaces

You can now configure the switch to display a login banner of up to 320 characters when an operator initiates a management session with the switch through any of the following methods:

- Telnet
- serial connection
- SSHv2 (SSHv1 does not include support for banners.)
- Web browser

In the factory default configuration, the switch displays the following default banner:

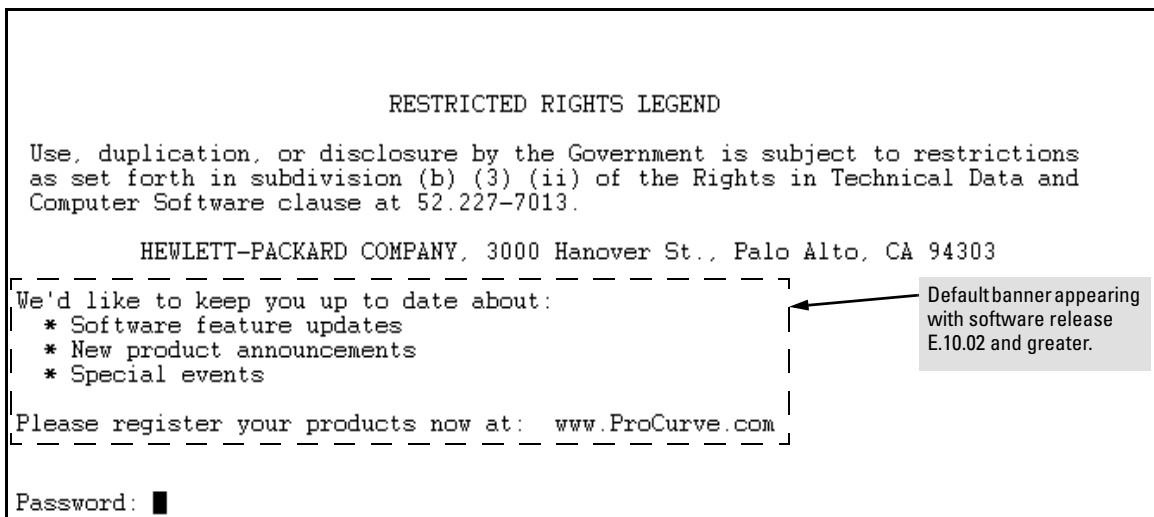


Figure 4-10. The Default Login Banner

Note

The switch's Web browser interface does not display the default banner.

Banner Operation with Telnet, Serial, or SSHv2 Access

When a system operator begins a login session, the switch displays the banner above the local password prompt or, if no password is configured, above the **Press any key to continue prompt**. Entering a correct password or, if no password is configured, pressing any key clears the banner from the CLI and displays the CLI prompt. (Refer to figure <zBlue> 4-10 on page 4-15.)

Banner Operation with Web Browser Access

When a system operator uses a Web browser to access the switch, the text of a non-default banner configured on the switch appears in a dedicated banner window with a link to the Web agent home page. Clicking on **To Home Page** clears the banner window and prompts the user for a password (if configured). Following entry of the correct username/password information (or if no username/password is required), the switch then displays either the Registration page or the switch's home page. Note that if the banner feature is disabled or if the switch is using the factory-default banner shown in figure <zBlue> 4-10, then the banner page does not appear in the Web browser when an operator initiates a login session with the switch.

Configuring and Displaying a Non-Default Banner

You can enable or disable banner operation using either the switch's CLI or an SNMP application. The steps include:

1. Enable non-default banner operation and define the endpoint delimiter for the banner.
2. Enter the desired banner text, including any specific line breaks you want.
3. Enter the endpoint delimiter.
4. Use **show banner motd** to display the current banner status.

Syntax: banner motd < *delimiter* >
 no banner motd

This command defines the single character used to terminate the banner text and enables banner text input. You can use any character except a blank space as a delimiter. The **no** form of the command disables the login banner feature.

< banner-text-string >

*The switch allows up to 320 banner characters, including blank spaces and CR-LF (Enter). (The tilde “~” and the delimiter defined by **banner motd <delimiter>** are not allowed as part of the banner text.) While entering banner text, you can backspace to edit the current line (that is, a line that has not been terminated by a CR-LF.) However, terminating a line in a banner by entering a CR-LF prevents any further editing of that line. To edit a line in a banner entry after terminating the line with a CR-LF requires entering the delimiter described above and then re-configuring new banner text.*

*The banner text string must terminate with the character defined by **banner motd <delimiter>**.*

Example of Configuring and Displaying a Banner

Suppose a system operator wanted to configure the following banner message on her company’s 5300xl switches:

```
This is a private system maintained by the  
Allied Widget Corporation.  
Unauthorized use of this system can result in  
civil and criminal penalties!
```

In this case, the operator will use the [Enter] key to create line breaks, blank spaces for line centering, and the % symbol to terminate the banner message.

```
ProCurve(config)# banner motd %  
Enter TEXT message. End with the character '%'  
    This is a private system maintained by the  
        Allied Widget Corporation.  
    Unauthorized use of this system can result in  
        civil and criminal penalties!%  
ProCurve(config)# write memory
```

Figure 4-11. Example of Configuring a Login Banner

To view the current banner configuration, use either the **show banner motd** or **show running** command.

```
ProCurve(config)# show banner motd

Banner Information

Banner status: Enabled
Configured Banner:

    This is a private system maintained by the
        Allied Widget Corporation.
    Unauthorized use of this system can result in
        civil and criminal penalties!
```

Figure 4-12. Example of show banner motd Output

```
ProCurve(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on release #E.10.02

hostname "ProCurve"
module 1 type J8161A
module 2 type J8161A
snmp-server community "notpublic" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B24
    ip address dhcp-bootp
    _exit_
banner motd "    This is a private system maintained by the
                Allied Widget Corporation.
    Unauthorized use of this system can result in
                civil and criminal penalties!"
password manager
password operator
```

Shows the current banner configuration.

Figure 4-13. The Current Banner Appears in the Switch's Running-Config File

The next time someone logs onto the switch's management CLI, the following appears:

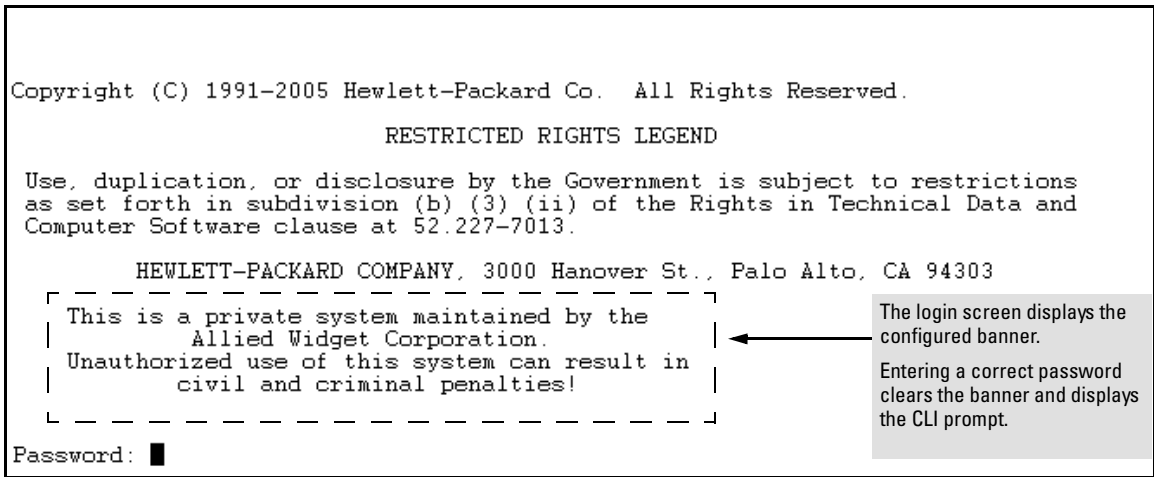


Figure 4-14. Example of CLI Result of the Login Banner Configuration

If someone uses a Web browser to log in to the switch interface, the following message appears:

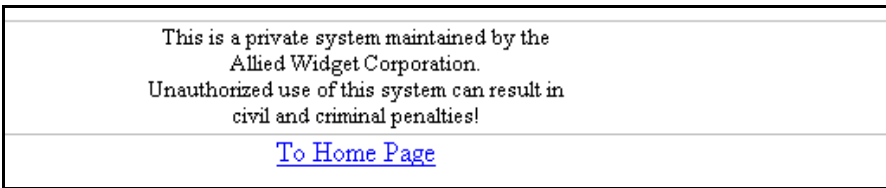


Figure 4-15. Example of Web Browser Interface Result of the Login Banner Configuration

Operating Notes

- The default banner appears only when the switch is in the factory default configuration. Using **no banner motd** deletes the currently configured banner text and blocks display of the default banner. The default banner is restored only if the switch is reset to its factory-default configuration.
- The switch supports one banner at any time. Configuring a new banner replaces any former banner configured on the switch.

- If the switch is configured with **ssh version 1** or **ssh version 1-or-2**, configuring the banner sets the SSH configuration to ssh version 2 and displays the following message in the CLI:

```
Warning: SSH version has been set to v2.
```





- If a banner is configured, the switch does not allow configuration with **ssh version 1** or **ssh version 1-or-2**. Attempting to do so produces the following error message in the CLI:

```
Banner has to be disabled first.
```

- If a banner is enabled on the switch, the Web browser interface displays the following link to the banner page:

Notice to all users

CLI Control and Editing

Keystrokes	Function
[Ctrl] [A]	Jumps to the first character of the command line.
[Ctrl] [B] or 	Moves the cursor back one character.
[Ctrl] [C]	Terminates a task and displays the command prompt.
[Ctrl] [D]	Deletes the character at the cursor.
[Ctrl] [E]	Jumps to the end of the current command line.
[Ctrl] [F] or 	Moves the cursor forward one character.
[Ctrl] [K]	Deletes from the cursor to the end of the command line.
[Ctrl] [L] or [Ctrl] [R]	Repeats current command line on a new line.
[Ctrl] [N] or 	Enters the next command line in the history buffer.
[Ctrl] [P] or 	Enters the previous command line in the history buffer.
[Ctrl] [U] or [Ctrl] [X]	Deletes from the cursor to the beginning of the command line.
[Ctrl] [W]	Deletes the last word typed.
[Esc] [B]	Moves the cursor backward one word.
[Esc] [D]	Deletes from the cursor to the end of the word.
[Esc] [F]	Moves the cursor forward one word.
[Backspace]	Deletes the first character to the left of the cursor in the command line.
[Spacebar]	Moves the cursor forward one character.

—This page is intentionally unused—

Using the Web Browser Interface

Contents

Overview	5-2
General Features	5-3
Starting an Web Browser Interface Session with the Switch	5-4
Using a Standalone Web Browser in a PC or UNIX Workstation	5-4
Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)	5-5
Tasks for Your First Web Browser Interface Session	5-7
Viewing the “First Time Install” Window	5-7
Security: Creating Usernames and Passwords in the Browser Interface	5-8
Entering a User Name and Password	5-10
Using a User Name	5-10
If You Lose the Password	5-10
Online Help for the Web Browser Interface	5-11
Support/Mgmt URLs Feature	5-12
Support URL	5-13
Help and the Management Server URL	5-13
Using the PCM Server for Switch Web HelpWeb Help	5-14
Status Reporting Features	5-16
The Overview Window	5-16
The Port Utilization and Status Displays	5-17
Port Utilization	5-17
Port Status	5-19
The Alert Log	5-20
Sorting the Alert Log Entries	5-20
Alert Types and Detailed Views	5-21
The Status Bar	5-22
Setting Fault Detection Policy	5-24

Overview

The web browser interface built into the switch lets you easily access the switch from a browser-based PC on your network. This lets you do the following:

- Optimize your network uptime by using the Alert Log and other diagnostic tools
- Make configuration changes to the switch
- Maintain security by configuring usernames and passwords

This chapter covers the following:

- General features (page 5-3).
- Starting a web browser interface session (page 5-4)
- Tasks for your first web browser interface session (page 5-7):
 - Creating usernames and passwords in the web browser interface (page 5-8)
 - Selecting the fault detection configuration for the Alert Log operation (page 5-24)
 - Getting access to online help for the web browser interface (page 5-11)
- Description of the web browser interface:
 - Overview window and tabs (page 5-16)
 - Port Utilization and Status displays (page 5-17)
 - Alert Log and Alert types (page 5-20)
 - Setting the Fault Detection Policy (page 5-24)

Note

You can disable access to the web browser interface by either executing **no web-management** at the Command Prompt or changing the **Web Agent Enabled** parameter setting to **No** (page 7-4).

General Features

The Web Browser Interface includes these features:

Switch Identity and Status:

- General system data
- Software version
- IP address
- Status Overview
- Port utilization
- Port counters
- Port status
- Alert log

Switch Configuration:

- Device view
- Port configuration
- VLAN configuration
- Fault detection
- Quality of service (QoS)
- Port monitoring (mirroring)
- System information
- IP configuration
- Support and management server URLs
- Device features (Spanning Tree On/Off, VLAN selection, and IGMP)
- Stacking (3400cl, 6400cl and 4200vl switches)

Switch Security:

- User names and passwords
- Authorized Addresses
- Intrusion Log
- SSL
- RADIUS authentication (Refer to the *Access Security Guide*.)

Switch Diagnostics:

- Ping/Link Test
- Device reset
- Configuration report

Starting an Web Browser Interface Session with the Switch

You can start a web browser session in the following ways:

- Using a standalone web browser on a network connection from a PC or UNIX workstation:
 - Directly connected to your network
 - Connected through remote access to your network
- Using a network management station running ProCurve Manager on your network

Using a Standalone Web Browser in a PC or UNIX Workstation

This procedure assumes that you are using a compatible web browser and that the switch is configured with an IP address accessible from your PC or workstation. (For more on assigning an IP address, refer to “IP Configuration” on page 8-2.)

1. Ensure that the Java™ applets are enabled for your browser. For more information on this topic, refer to your browser’s online Help.
2. Use the web browser to access the switch. If your network includes a Domain Name Server (DNS), your switch’s IP address may have a name associated with it (for example, **switch5308**) that you can type in the **Location or Address** field instead of the IP address. Using DNS names typically improves browser performance. Contact your network administrator to enquire about DNS names associated with your ProCurve switch.

Type the IP address (or DNS name) of the switch in the browser **Location or Address** (URL) field and press **[Enter]**. (It is not necessary to include **http://**.)

switch5308 **[Enter]** (example of a DNS-type name)

10.11.12.195 **[Enter]** (example of an IP address)

Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)

ProCurve Manager and ProCurve Manager Plus are designed for installation on a network management workstation. For this reason, the system requirements are different from the system requirements for accessing the switch's web browser interface from a non-management PC or workstation. For PCM and PCM+ requirements, refer to the information provided with the software.

This procedure assumes that:

- You have installed the recommended web browser on a PC or workstation that serves as your network management station.
- The networked device you want to access has been assigned an IP address and (optionally) a DNS name, and has been discovered by PCM or PCM+. (For more on assigning an IP address, refer to “IP Configuration” on page 8-2.)

To establish a web browser session with PCM or PCM+ running, do the following on the network management station:

1. Make sure the Java™ applets are enabled for your web browser. If they are not, refer to the web browser online Help for specific information on enabling the Java applets.
2. In the **Interconnected Devices** listing under **Network Manager Home** (in the PCM/PCM+ sidebar), right-click on the model number of the device you want to access.
3. The web browser interface automatically starts with the Status Overview window displayed for the selected device, as shown in figure 5-1.

Note

If the Registration window appears, click on the **Status** tab.

Using the Web Browser Interface

Starting an Web Browser Interface Session with the Switch

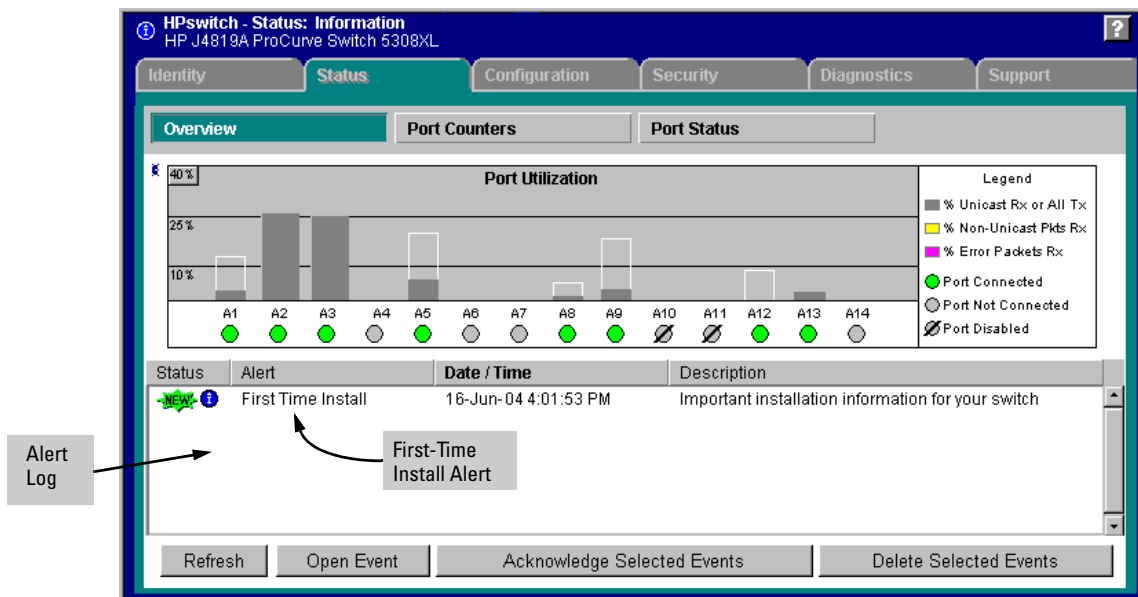


Figure 5-1. Example of Status Overview Screen

Tasks for Your First Web Browser Interface Session

The first time you access the web browser interface, there are three tasks you should perform:

- Review the “First Time Install” window
- Set Manager and Operator passwords
- Set access to the web browser interface online help

Viewing the “First Time Install” Window

When you access the switch’s web browser interface for the first time, the Alert log contains a “First Time Install” alert, as shown in figure 5-2. This gives you information about first time installations, and provides an immediate opportunity to set passwords for security and to specify a Fault Detection policy, which determines the types of messages that will be displayed in the Alert Log.

Double click on **First Time Install** in the Alert log (figure 5-1 on page 5-6). The web browser interface then displays the “First Time Install” window, below.

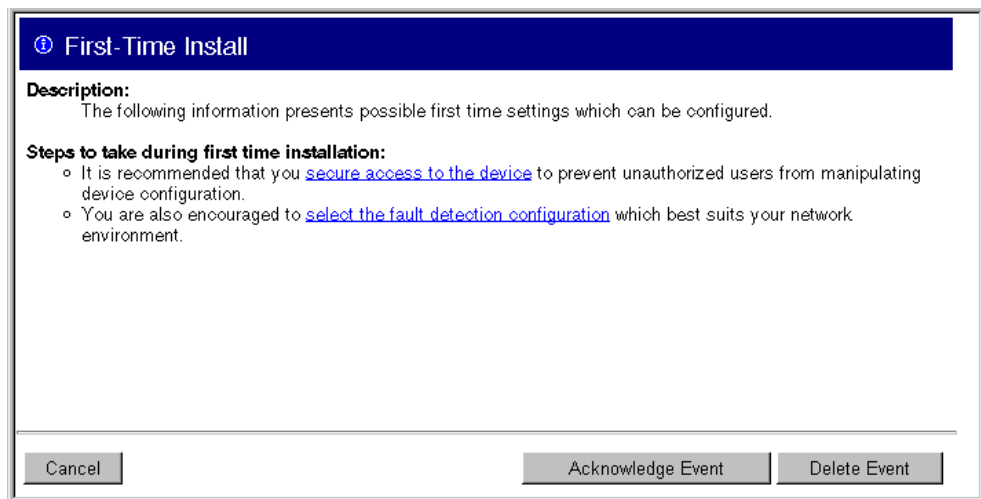


Figure 5-2. First-Time Install Window

This window is the launching point for the basic configuration you need to perform to set web browser interface passwords for maintaining security and a fault detection policy, which determines the types of messages that the Alert Log displays.

To set web browser interface passwords, click on **secure access to the device** to display the Device Passwords screen, and then go to the next page. (You can also access the password screen by clicking on the **Security** tab.)

To set Fault Detection policy, click on **select the fault detection configuration** in the second bullet in the window and go to the section, “Setting Fault Detection Policy” on page 5-24. (You can also access the password screen by clicking on the **Configuration** tab, and then the **[Fault Detection]** key.)

Security: Creating Usernames and Passwords in the Browser Interface

Note

On 5300xl switches running software release E.09.xx, you can also configure RADIUS authentication for web browser interface access. For more information, refer to the chapter titled “RADIUS Authentication and Accounting” in the *Access Security Guide* for your switch.

You may want to create both a username and a password to create access security for your switch. There are two levels of access to the interface that can be controlled by setting user names and passwords:

- **Operator Setting.** An Operator-level user name and password allows read-only access to most of the web browser interface, but prevents access to the Security window.
- **Manager Setting.** A Manager-level user name and password allows full read/write access to the web browser interface.

Switch-1 - Status: Information
HP J4850A ProCurve Switch 5304XL

Identity Status Configuration **Security** Diagnostics Support

Device Passwords Authorized Addresses Port Security Intrusion Log SSL

Read-Only Access

Operator User Name:
Operator Password:
Confirm Operator Password:

Read-Write Access

Manager User Name:
Manager Password:
Confirm Manager Password:

Apply Changes Clear Changes

Figure 5-3. The Device Passwords Window

To set the passwords:

1. Access the Device Passwords screen by one of the following methods:
 - If the Alert Log includes a “First Time Install” event entry, double click on this event, then, in the resulting display, click on the **secure access to the device** link.
 - Select the **Security** tab.
2. Click in the appropriate box in the **Device Passwords** window and enter user names and passwords. You will be required to repeat the password strings in the confirmation boxes.

Both the user names and passwords can be up to 16 printable ASCII characters.

3. Click on **[Apply Changes]** to activate the user names and passwords.

Note

Passwords you assign in the web browser interface will overwrite previous passwords assigned in either the web browser interface, the CLI, or the menu interface. That is, the most recently assigned passwords are the switch's passwords, regardless of which interface was used to assign the string.

Entering a User Name and Password

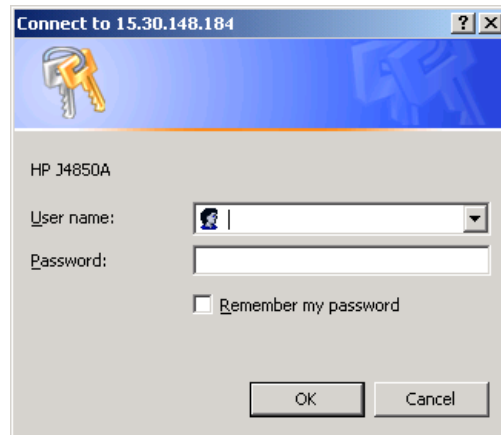


Figure 5-4. Example of the Password Prompt in the Web Browser Interface

The manager and operator passwords are used to control access to all switch interfaces. Once set, you will be prompted to supply the password every time you try to access the switch through any of its interfaces. The password you enter determines the capability you have during that session:

- Entering the manager password gives you full read/write/troubleshooting capabilities
- Entering the operator password gives you read and limited troubleshooting capabilities.

Using a User Name

If you also set user names in the web browser interface screen, you must supply the correct user name for web browser interface access. If a user name has not been set, then leave the User Name field in the password window blank.

Note that the Command Prompt and switch console interfaces use only the password, and do not prompt you for the User Name.

If You Lose the Password

If you lose the passwords, you can clear them by pressing the Clear button on the front of the switch. *This action deletes all password and user name protection from all of the switch's interfaces.*

The Clear button is provided for your convenience, but its presence means that if you are concerned with the security of the switch configuration and operation, you should make sure the switch is installed in a secure location, such as a locked wiring closet. (For more information, refer to “Front Panel Security” in the chapter titled “Configuring Username and Password Security” in the Access Security Guide for your switch.)

Online Help for the Web Browser Interface

Online Help is available for the web browser interface. You can use it by clicking on the question mark button in the upper right corner of any of the web browser interface screens.

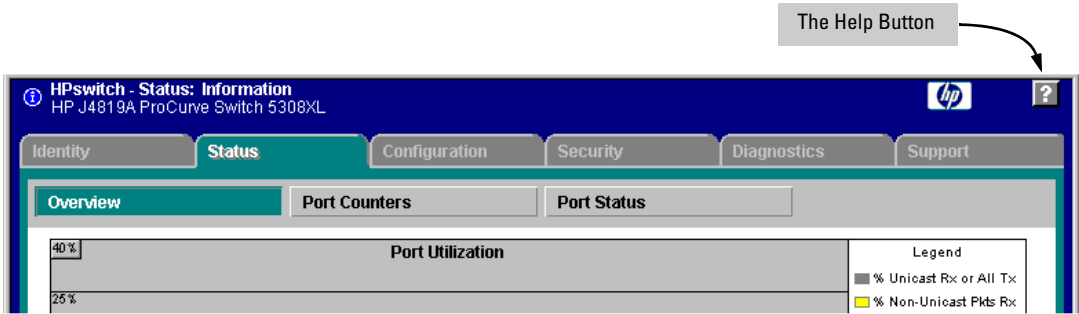


Figure 5-5. The Help Button

Context-sensitive help is provided for the screen you are on.

Note

To access the online Help for the web browser interface, you need either ProCurve Manager (version 1.5 or greater) installed on your network or an active connection to the World Wide Web. Otherwise, Online help for the web browser interface will not be available.

For more on Help access and operation, refer to “Help and the Management Server URL” on page 5-13.

Support/Mgmt URLs Feature

The Support/Mgmt URLs window enables you to change the World Wide Web Universal Resource Locator (URL) for two functions:

- **Support URL** – A support information site for your switch
- **Management Server URL** – The web site for web browser online Help

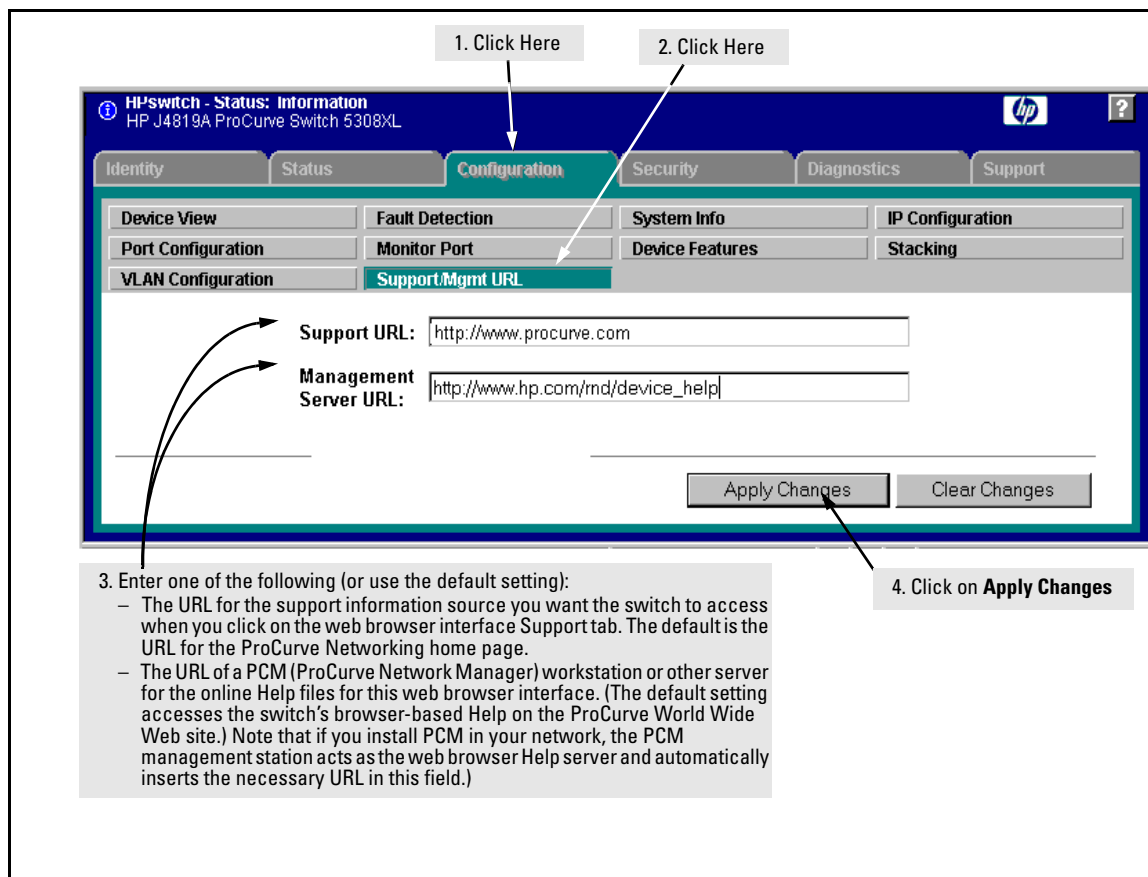


Figure 5-6. The Default Support/Mgmt URLs Window

Support URL

This is the site the switch accesses when you click on the **Support** tab on the web browser interface. The default URL is:

www.procurve.com

which is the World Wide Web site for ProCurve networking products. Click on **technical support** on that page to get support information regarding your switch, including white papers, software updates, and more.

As an alternative, you can replace the ProCurve URL with the URL for a local site used for logging reports on network performance or other support activities.

Help and the Management Server URL

The **Management Server URL** field specifies the URL the switch uses to find online Help for the web browser interface.

- If you install PCM (ProCurve Manager) in your network, the PCM management station acts as the web browser Help server for the switch and automatically inserts the necessary URL in this field.)
- In the default configuration (and if PCM is not running on your network) this field is set to the URL for accessing online Help from the ProCurve Networking web site:

www.hp.com/rnd/device_help

Using this option, the Help files are automatically available if your workstation can access the World Wide Web. In this case, if Online Help fails to operate, ensure that the above URL appears in the **Management Server URL** field shown in figure 5-7:

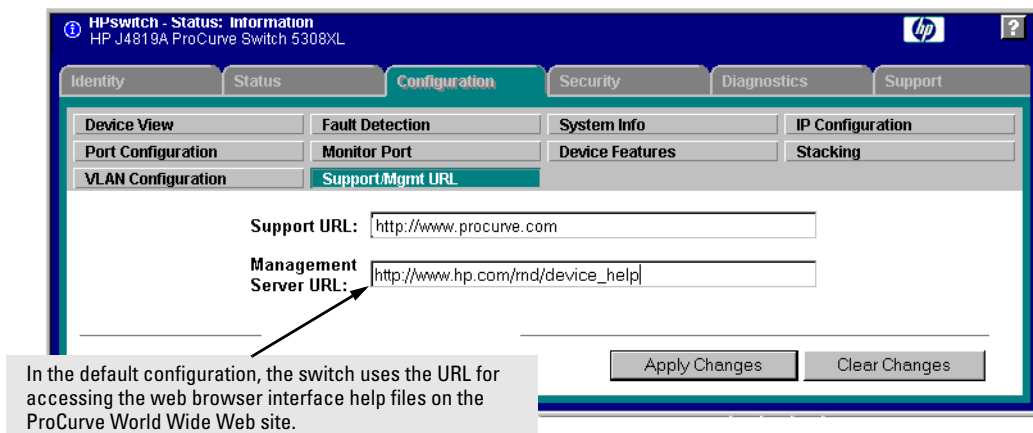


Figure 5-7. How To Access Web Browser Interface Online Help

Using the PCM Server for Switch Web Help

For ProCurve devices that support the “Web Help” feature, you can use the PCM server to host the switch help files for devices that do not have HTTP access to the ProCurve Support Web site.

1. Go to the ProCurve Support web site to get the Device Help files:
www.hp.com/rnd/device_help/
2. Copy the Web help files to the PCM server, under:
C:\\program files\\hewlett-packard\\pnm\\server\\webroot\\
rnd\\sevice_help\\help\\hpwnd\\webhelp
3. Add an entry, or edit the existing entry in the Discovery portion of the global properties (globalprops.prp) in PCM to redirect the switches to the help files on the PCM server. For example:

```
Global {  
TempDir=data/temp  
...  
Discovery{  
...  
...  
DeviceHelpUrlRedirect=http://15.29.37.12.8040/rnd/device_help  
...  
}
```

}

You will enter the IP address for your PCM server. 8040 is the standard port number to use.

4. Restart the Discovery process for the change to be applied.

Note

Changing the Discovery's Global properties file will redirect the Device Help URL for all devices.

If you just want to change the Device Help URL for a particular device, then go to the Configuration tab on the Web UI for that device and select the "Support/Mgmt URL" button. Edit the entry in the "Management Server URL" field for the device to point to the PCM server; for example:

http://15.29.37.12.8040/rnd/device_help

Status Reporting Features

Browser elements covered in this section include:

- The Overview window (below)
- Port utilization and status (page 5-17)
- The Alert log (page 5-20)
- The Status bar (page 5-22)

The Overview Window

The Overview Window is the home screen for any entry into the web browser interface. The following figure identifies the various parts of the screen.

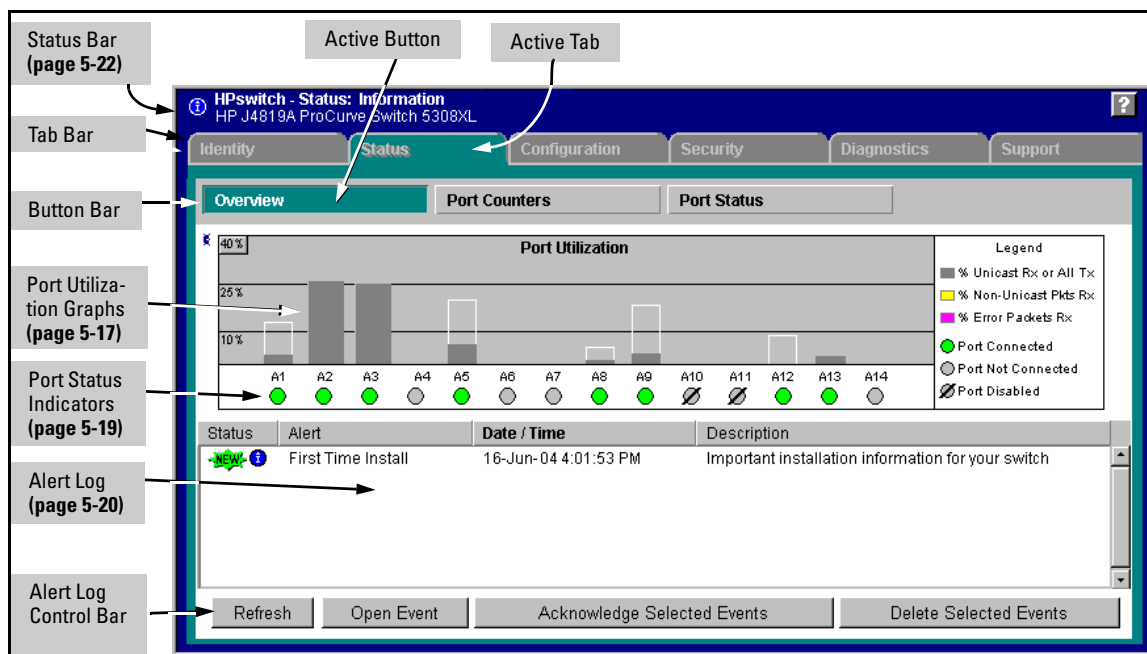


Figure 5-8. The Status Overview Window

Policy Management and Configuration. PCM can perform network-wide policy management and configuration of your switch. The Management Server URL field (page 5-13) shows the URL for the management station performing that function. For more information, refer to the documentation provided with the PCM software.

The Port Utilization and Status Displays

The Port Utilization and Status displays show an overview of the status of the switch and the amount of network activity on each port. The following figure shows a sample reading of the Port Utilization and Port Status.

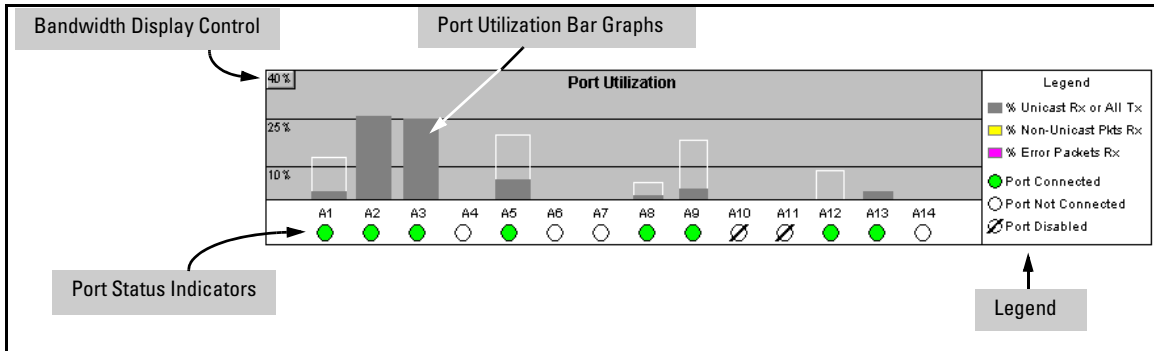


Figure 5-9. The Graphs Area

Port Utilization

The Port Utilization bar graphs show the network traffic on the port with a breakdown of the packet types that have been detected (unicast packets, non-unicast packets, and error packets). The Legend identifies traffic types and their associated colors on the bar graph:

- **% Unicast Rx & All Tx:** This is all unicast traffic received and all transmitted traffic of any type. This indicator (a blue color on many systems) can signify either transmitted or received traffic.
- **% Non-Unicast Pkts Rx:** All multicast and broadcast traffic received by the port. This indicator (a gold color on many systems) enables you to know “at-a-glance” the source of any non-unicast traffic that is causing high utilization of the switch. For example, if one port is receiving heavy broadcast or multicast traffic, all ports will become highly utilized. By color-coding the received broadcast and multicast utilization, the bar graph quickly and easily identifies the offending port. This makes it faster and easier to discover the exact source of the heavy traffic because you don’t have to examine port counter data from several ports.

- **% Error Pkts Rx:** All error packets received by the port. (This indicator is a reddish color on many systems.) Although errors received on a port are not propagated to the rest of the network, a consistently high number of errors on a specific port may indicate a problem on the device or network segment connected to the indicated port.
- **Maximum Activity Indicator:** As the bars in the graph area change height to reflect the level of network activity on the corresponding port, they leave an outline to identify the maximum activity level that has been observed on the port.

Utilization Guideline. A network utilization of 40% is considered the maximum that a typical Ethernet-type network can experience before encountering performance difficulties. If you observe utilization that is consistently higher than 40% on any port, click on the Port Counters button to get a detailed set of counters for the port.

To change the amount of bandwidth the Port Utilization bar graph shows. Click on the bandwidth display control button in the upper left corner of the graph. (The button shows the current scale setting, such as 40%.) In the resulting menu, select the bandwidth scale you want the graph to show (3%, 10%, 25%, 40%, 75%, or 100%), as shown in figure figure 5-10.

Note that when viewing activity on a gigabit port, you may want to select a lower value (such as 3% or 10%). This is because the bandwidth utilization of current network applications on gigabit links is typically minimal, and may not appear on the graph if the scale is set to show high bandwidth utilization.

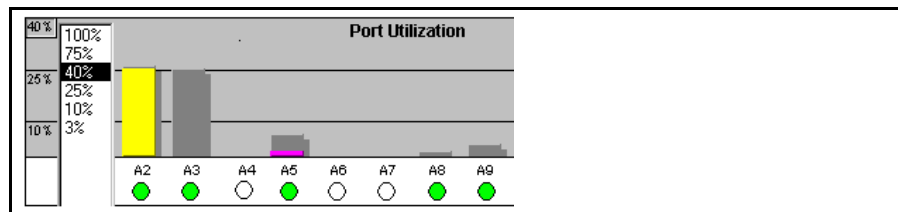


Figure 5-10. Changing the Graph Area Scale

To display values for each graph bar. Hold the mouse cursor over any of the bars in the graph, and a pop-up display is activated showing the port identification and numerical values for each of the sections of the bar, as shown in figure 5-11 (next).

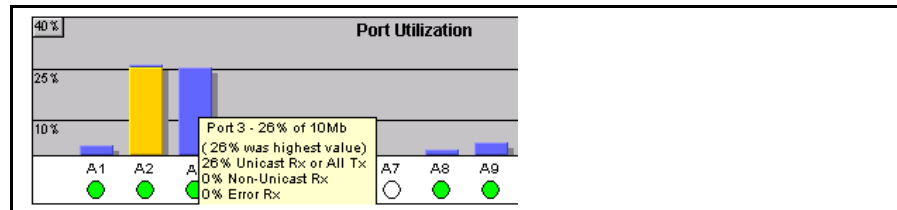


Figure 5-11. Display of Numerical Values for the Bar

Port Status

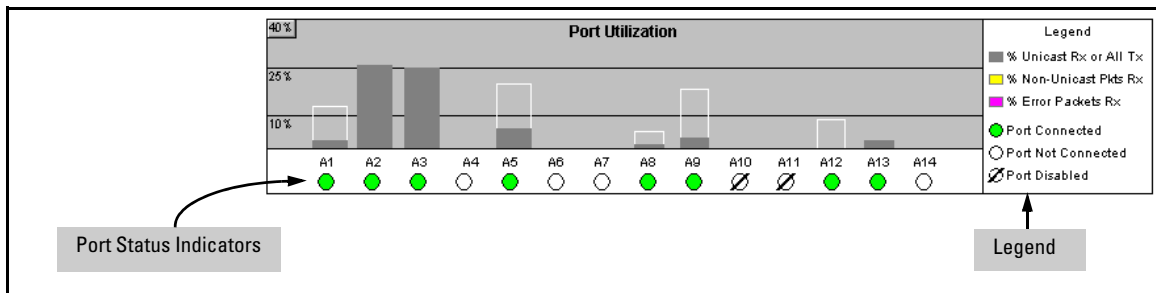


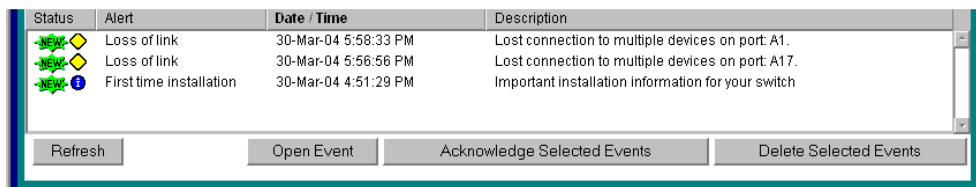
Figure 5-12. The Port Status Indicators and Legend







The Port Status indicators show a symbol for each port that indicates the general status of the port. There are four possible statuses:

- **Port Connected** – the port is enabled and is properly connected to an active network device.
- **Port Not Connected** – the port is enabled but is not connected to an active network device. A cable may not be connected to the port, or the device at the other end may be powered off or inoperable, or the cable or connected device could be faulty.
- **Port Disabled** – the port has been configured as disabled through the web browser interface, the switch console, or SNMP network management.
- **Port Fault-Disabled** – a fault condition has occurred on the port that has caused it to be auto-disabled. Note that the Port Fault-Disabled symbol will be displayed in the legend only if one or more of the ports is in that status. See appendix B, "Monitoring and Analyzing Switch Operation" for more information.

The Alert Log

The web browser interface Alert Log, shown in the lower half of the screen, shows a list of network occurrences, or *alerts*, that were detected by the switch. Typical alerts are **Broadcast Storm**, indicating an excessive number of broadcasts received on a port, and **Problem Cable**, indicating a faulty cable. A full list of alerts is shown in the table on page 5-21.



Status	Alert	Date / Time	Description
 	Loss of link	30-Mar-04 5:58:33 PM	Lost connection to multiple devices on port: A1.
 	Loss of link	30-Mar-04 5:56:56 PM	Lost connection to multiple devices on port: A17.
 	First time installation	30-Mar-04 4:51:29 PM	Important installation information for your switch

Refresh Open Event Acknowledge Selected Events Delete Selected Events

Figure 5-13. Example of the Alert Log

Each alert has the following fields of information:

- **Status** – The level of severity of the event generated. Severity levels can be Information, Normal, Warning, and Critical. If the alert is new (has not yet been acknowledged), the New symbol is also in the Status column.
- **Alert** – The specific event identification.
- **Date/Time** – The date and time the event was received by the web browser interface. This value is shown in the format: **DD-MM-YY HH:MM:SS AM/PM**, for example, **16-Sep-99 7:58:44 AM**.
- **Description** – A short narrative statement that describes the event. For example, **Excessive CRC/Alignment errors on port: 8**.

Sorting the Alert Log Entries

The alerts are sorted, by default, by the Date/Time field with the most recent alert listed at the top of the list. The second most recent alert is displayed below the top alert and so on. If alerts occurred at the same time, the simultaneous alerts are sorted by order in which they appear in the MIB.

Bold characters in a column heading indicate that the alert field alert log entries. You can sort by any of the other columns by clicking on the column heading. The **Alert** and **Description** columns are sorted alphabetically, while the **Status** column is sorted by severity type, with more critical severity indicators appearing above less critical indicators.

Alert Types and Detailed Views

As of April, 2004, the web browser interface generates the following alert types:

- Auto Partition
- Backup Transition
- Excessive broadcasts
- Excessive CRC/alignment errors
- Excessive jabbering
- Excessive late collisions
- First Time Install
- Full-Duplex Mismatch
- Half-Duplex Mismatch
- High collision or drop rate
- Loss of Link
- Mis-Configured SQE
- Network Loop
- Polarity Reversal
- Security Violation
- Stuck 10BaseT Port
- Too many undersized (runt)/giant packets
- Transceiver Hot Swap

Note

When troubleshooting the sources of alerts, it may be helpful to check the switch's Port Status and Port Counter windows, or use the CLI or menu interface to view the switch's Event Log.

When you double click on an Alert Entry, the web browser interface displays a separate window showing information about the event. This view includes a description of the problem and a possible solution. It also provides three management buttons:

- **Acknowledge Event** – removes the New symbol from the log entry
- **Delete Event** – removes the alert from the Alert Log
- **Cancel** – closes the detail view with no change to the status of the alert and returns you to the Overview screen.

For example, figure 5-14 shows a sample detail view describing an Excessive CRC/Alignment Error alert.

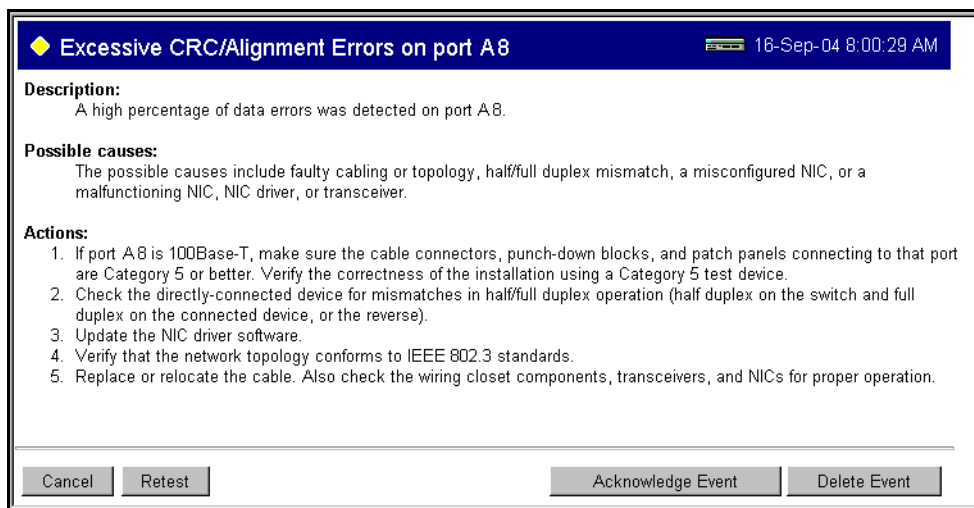


Figure 5-14. Example of Alert Log Detail View

The Status Bar

The Status Bar appears in the upper left corner of the web browser interface window. Figure 5-15 shows an expanded view of the status bar.

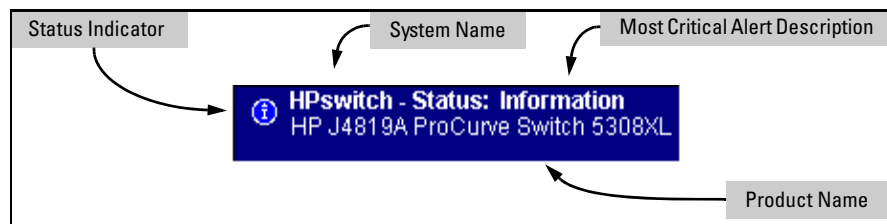






Figure 5-15. Example of the Status Bar

The Status bar includes four objects:

- **Status Indicator.** Indicates, by icon, the severity of the most critical alert in the current display of the Alert Log. This indicator can be one of four shapes and colors, as shown below.

Table 5-1. Status Indicator Key

Color	Switch Status	Status Indicator Shape
Blue	Normal Activity; "First time installation" information available in the Alert log.	
Green	Normal Activity	
Yellow	Warning	
Red	Critical	

- **System Name.** The name you can configure for the switch by using the **System Info** window (under the **Configuration** tab), the **hostname < ascii-string >** command in the CLI, or the **System Name** field in the "System Information" screen in the System Info screen of the menu interface.
- **Most Critical Alert Description.** A brief description of the earliest, unacknowledged alert with the current highest severity in the Alert Log, appearing in the right portion of the Status Bar. In instances where multiple critical alerts have the same severity level, only the earliest unacknowledged alert is deployed in the Status bar.
- **Product Name.** The product name of the switch to which you are connected in the current web browser interface session.

Setting Fault Detection Policy

One of the powerful features in the web browser interface is the Fault Detection facility. For your switch, this feature controls the types of alerts reported to the Alert Log based on their level of severity.

Set this policy in the Fault Detection window (figure 5-16).

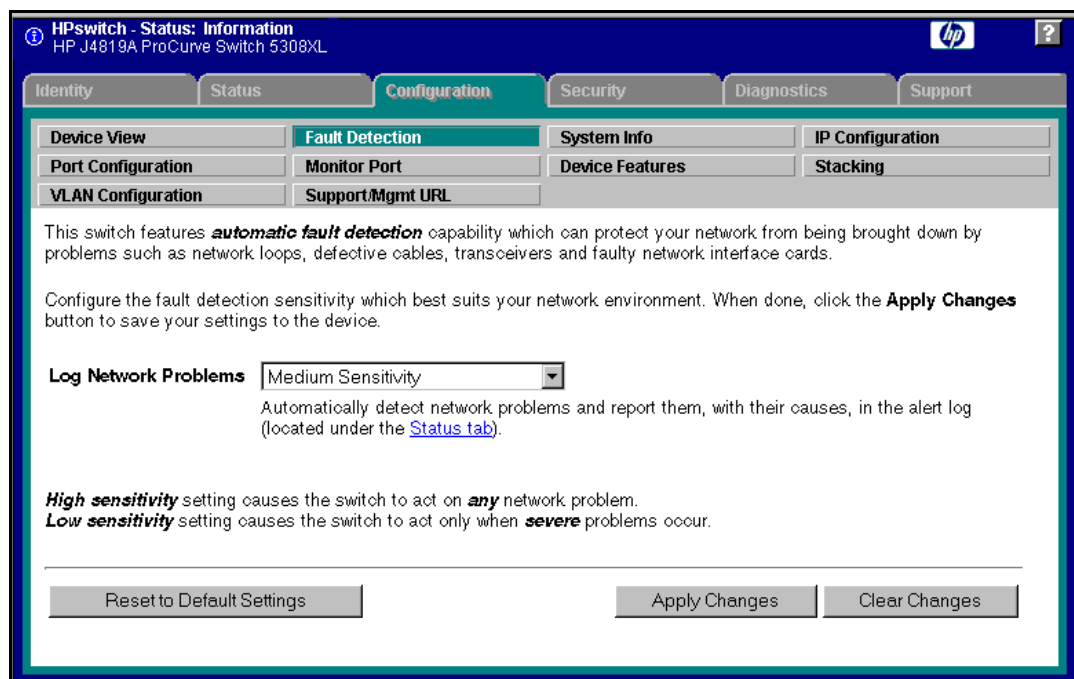


Figure 5-16. The Fault Detection Window

The Fault Detection screen contains a list box for setting fault detection and response policy, and enables you to set the sensitivity level at which a network problem should generate an alert and send it to the Alert Log.

To provide the most information on network problems in the Alert Log, the recommended sensitivity level for **Log Network Problems** is **High Sensitivity**. The Fault Detection settings are:

- **High Sensitivity.** This policy directs the switch to send all alerts to the Alert Log. This setting is most effective on networks that have none or few problems.
- **Medium Sensitivity.** This policy directs the switch to send alerts related to network problems to the Alert Log. If you want to be notified of problems which cause a noticeable slowdown on the network, use this setting.
- **Low Sensitivity.** This policy directs the switch to send only the most severe alerts to the Alert Log. This policy is most effective on a network where there are normally a lot of problems and you want to be informed of only the most severe ones.
- **Never.** Disables the Alert Log and transmission of alerts (traps) to the management server (in cases where a network management tool such as ProCurve Manager is in use). Use this option when you don't want to use the Alert Log.

The Fault Detection Window also contains three Change Control Buttons:

- **Apply Changes.** This button stores the settings you have selected for all future sessions with the web browser interface until you decide to change them.
- **Clear Changes.** This button removes your settings and returns the settings for the list box to the level it was at in the last saved detection-setting session.
- **Reset to Default Settings.** This button reverts the policy setting to Medium Sensitivity for Log Network Problems.

—This page left blank intentionally—

Switch Memory and Configuration

Contents

Overview	6-3
Overview of Configuration File Management	6-3
Using the CLI To Implement Configuration Changes	6-6
Using the Menu and Web Browser Interfaces To Implement Configuration Changes	6-9
Menu: Implementing Configuration Changes	6-9
Using Save and Cancel in the Menu Interface	6-10
Rebooting from the Menu Interface	6-11
Web: Implementing Configuration Changes	6-12
Using Primary and Secondary Flash Image Options	6-13
Displaying the Current Flash Image Data	6-13
Switch Software Downloads	6-15
Local Switch Software Replacement and Removal	6-16
Rebooting the Switch	6-18
Operating Notes	6-21
Multiple Configuration Files on 5300xl and 4200vl Switches	6-22
General Operation	6-24
Transitioning to Multiple Configuration Files	6-26
Listing and Displaying Startup-Config Files	6-27
Viewing the Startup-Config File Status with Multiple Configuration Enabled	6-27
Displaying the Content of A Specific Startup-Config File	6-29
Changing or Overriding the Reboot Configuration Policy	6-29
Managing Startup-Config Files in the Switch	6-31
Renaming an Existing Startup-Config File	6-32
Creating a New Startup-Config File	6-32
Erasing a Startup-Config File	6-34

Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration	6-35
Transferring Startup-Config Files To or From a Remote Server	6-37
TFTP: Copying a Configuration File to a Remote Host	6-37
TFTP: Copying a Configuration File from a Remote Host	6-37
Xmodem: Copying a Configuration File to a Serially Connected Host	6-38
Xmodem: Copying a Configuration from a Serially Connected Host	6-38
Operating Notes for Multiple Configuration Files	6-39

Overview

This chapter describes:

- How switch memory manages configuration changes
 - How the CLI implements configuration changes
 - How the menu interface and web browser interface implement configuration changes
 - How the switch provides software options through primary/secondary flash images
 - How to use the switch's primary and secondary flash options, including displaying flash information, booting or restarting the switch, and other topics
-

Overview of Configuration File Management

The switch maintains two configuration files, the *running-config* file and the *startup-config* file.

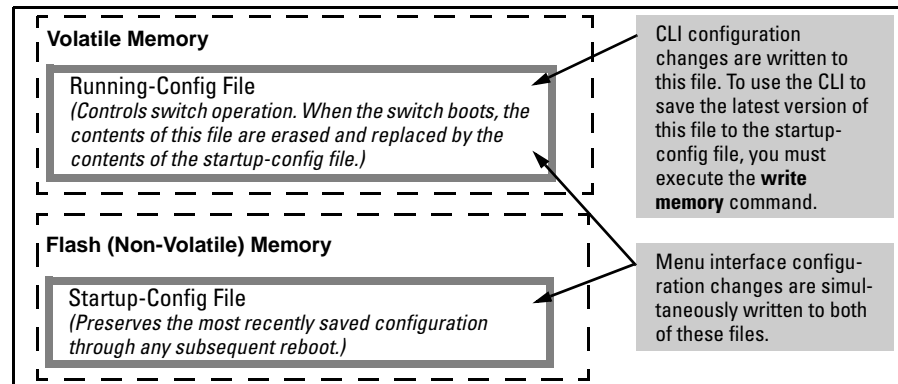


Figure 6-1. Conceptual Illustration of Switch Memory Operation

- **Running Config File:** Exists in volatile memory and controls switch operation. If no configuration changes have been made in the CLI since the switch was last booted, the running-config file is identical to the startup-config file.
- **Startup-config File:** Exists in flash (non-volatile) memory and is used to preserve the most recently-saved configuration as the “permanent” configuration.

Booting the switch replaces the current running-config file with a new running-config file that is an exact copy of the current startup-config file.

Note

Any of the following actions boots the switch:

- Executing the **boot** or the **reload** command in the CLI
- Executing the **boot** command in the menu interface
- Pressing the Reset button on the front of the switch
- Removing, then restoring power to the switch

For more on reboots and the switch’s dual-flash images, see “Using Primary and Secondary Flash Image Options” on page 6-13.

Options for Saving a New Configuration. Making one or more changes to the running-config file creates a new operating configuration. *Saving* a new configuration means to overwrite (replace) the current startup-config file with the current running-config file. This means that if the switch subsequently reboots for any reason, it will resume operation using the new configuration instead of the configuration previously defined in the startup-config file. There are three ways to save a new configuration:

- **In the CLI:** Use the **write memory** command. This overwrites the current startup-config file with the contents of the current running-config file.
- **In the menu interface:** Use the **Save** command. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the menu interface screen.
- **In the web browser interface:** Use the **[Apply Changes]** button or other appropriate button. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the web browser interface window.

Note that using the CLI instead of the menu or web browser interface gives you the option of changing the running configuration without affecting the startup configuration. This allows you to test the change without making it

“permanent”. When you are satisfied that the change is satisfactory, you can make it permanent by executing the **write memory** command. For example, suppose you use the following command to disable port 5:

```
ProCurve(config)# interface ethernet 5 disable
```

The above command disables port 5 in the running-config file, but not in the startup-config file. Port 5 remains disabled only until the switch reboots. If you want port 5 to remain disabled through the next reboot, use **write memory** to save the current running-config file to the startup-config file in flash memory.

```
ProCurve(config)# write memory
```

If you use the CLI to make a configuration change and then change from the CLI to the Menu interface without first using write memory to save the change to the startup-config file, then the switch prompts you to save the change. For example, if you use the CLI to create VLAN 20, and then select the menu interface, VLAN 20 is configured in the running-config file, but not in the startup-config file. In this case you will see:

```
ProCurve(config)# vlan 20
ProCurve(config)# menu
Do you want to save current configuration [y/n]?
```

If you type **[Y]**, the switch overwrites the startup-config file with the running-config file, and your configuration change(s) will be preserved across reboots. If you type **[N]**, your configuration change(s) will remain only in the running-config file. In this case, if you do not subsequently save the running-config file, your unsaved configuration changes will be lost if the switch reboots for any reason.

Storing and Retrieving Configuration Files. You can store or retrieve a backup copy of the startup-config file on another device. For more information, see appendix A, “Transferring an Operating System or Startup-Config File”

Using the CLI To Implement Configuration Changes

The CLI offers these capabilities:

- Access to the full set of switch configuration features
- The option of testing configuration changes before making them permanent

How To Use the CLI To View the Current Configuration Files. Use **show** commands to view the configuration for individual features, such as port status or Spanning Tree Protocol. However, to view either the entire startup-config file or the entire running-config file, use the following commands:

- **show config** — Displays a listing of the current startup-config file.
- **show running-config** — Displays a listing of the current running-config file.
- **write terminal** — Displays a listing of the current running-config file.
- **show config status** — Compares the startup-config file to the running-config file and lists one of the following results:
 - If the two configurations are the same you will see:
 - Running configuration is the same as the startup configuration.
 - If the two configurations are different, you will see:
 - Running configuration has been changed and needs to be saved.

Note

Show config, **show running-config**, and **write terminal** commands display the configuration settings that differ from the switch's factory-default configuration.

How To Use the CLI To Reconfigure Switch Features. Use this procedure to permanently change the switch configuration (that is, to enter a change in the startup-config file).

1. Use the appropriate CLI commands to reconfigure the desired switch parameters. This updates the selected parameters in the running-config file.
2. Use the appropriate **show** commands to verify that you have correctly made the desired changes.

3. Observe the switch's performance with the new parameter settings to verify the effect of your changes.
4. When you are satisfied that you have the correct parameter settings, use the **write memory** command to copy the changes to the startup-config file.

Syntax: write memory

For example, the default port mode setting is **auto**. Suppose that your network uses Cat 3 wiring and you want to connect the switch to another autosensing device capable of 100 Mbps operation. Because 100 Mbps over Cat 3 wiring can introduce transmission problems, the recommended port mode is **auto-10**, which allows the port to negotiate full- or half-duplex, but restricts speed to 10 Mbps. The following command configures port A5 to auto-10 mode in the running-config file, allowing you to observe performance on the link without making the mode change permanent.

```
ProCurve(config)# interface e a5 speed-duplex auto-10
```

After you are satisfied that the link is operating properly, you can save the change to the switch's permanent configuration (the startup-config file) by executing the following command:

```
ProCurve(config)# write memory
```

The new mode (**auto-10**) on port A5 is now saved in the startup-config file, and the startup-config and running-config files are identical. If you subsequently reboot the switch, the **auto-10** mode configuration on port A5 will remain because it is included in the startup-config file.

How To Cancel Changes You Have Made to the Running-Config File.

If you use the CLI to change parameter settings in the running-config file, and then decide that you don't want those changes to remain, you can use either of the following methods to remove them:

- Manually enter the earlier values you had for the changed settings. (This is recommended if you want to restore a small number of parameter settings to their previous boot-up values.)
- Update the running-config file to match the startup-config file by rebooting the switch. (This is recommended if you want to restore a larger number of parameter settings to their previous boot-up values.)

If you use the CLI to change a parameter setting, and then execute the **boot** command without first executing the **write memory** command to save the change, the switch prompts you to specify whether to save the changes in the current running-config file. For example:

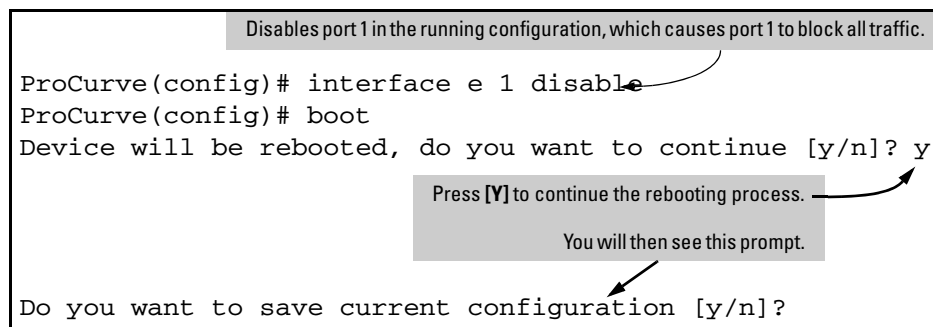


Figure 6-2. Boot Prompt for an Unsaved Configuration

The above prompt means that one or more parameter settings in the running-config file differ from their counterparts in the startup-config file and you need to choose which config file to retain and which to discard.

- If you want to update the startup-config file to match the running-config file, press **[Y]** for “yes”. (This means that the changes you entered in the running-config file will be saved in the startup-config file.)
- If you want to discard the changes you made to the running-config file so that it will match the startup-config file, then press **[N]** for “no”. (This means that the switch will discard the changes you entered in the running-config file and will update the running-config file to match the startup-config file.)

Note

If you use the CLI to make a change to the running-config file, you should either use the **write memory** command or select the save option allowed during a reboot (figure 6-2, above) to save the change to the startup-config file. That is, if you use the CLI to change a parameter setting, but then reboot the switch from either the CLI or the menu interface without first executing the **write memory** command in the CLI, the current startup-config file will replace the running-config file, and any changes in the running-config file will be lost.

Using the **Save** command in the menu interface does not save a change made to the running config by the CLI unless you have also made a configuration change in the menu interface. Also, the menu interface displays the current running-config values. Thus, where a parameter setting is accessible from both the CLI and the menu interface, if you change the setting in the CLI, the new value will appear in the menu interface display for that parameter. *However, as indicated above, unless you also make a configuration change in the menu interface, only the **write memory** command in the CLI will actually save the change to the startup-config file.*

How To Reset the startup-config and running-config Files to the Factory Default Configuration. This command reboots the switch, replacing the contents of the current startup-config and running-config files with the factory-default startup configuration.

Syntax: erase startup-config

For example:

```
ProCurve(config)# erase startup-config
Configuration will be deleted and device rebooted, continue [y/n]?
```

Press **[Y]** to replace the current configuration with the factory default configuration and reboot the switch. Press **[N]** to retain the current configuration and prevent a reboot.

Using the Menu and Web Browser Interfaces To Implement Configuration Changes

The menu and web browser interfaces offer these advantages:

- Quick, easy menu or window access to a subset of switch configuration features
- Viewing several related configuration parameters in the same screen, with their default and current settings
- Immediately changing both the running-config file and the startup-config file with a single command

Menu: Implementing Configuration Changes

You can use the menu interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change in the menu interface, you simultaneously change both the running-config file and the startup-config file.

Note

The only exception to this operation are two VLAN-related parameter changes that require a reboot—described under “Rebooting To Activate Configuration Changes” on page 6-11.

Using **Save** and **Cancel** in the Menu Interface

For any configuration screen in the menu interface, the Save command:

1. Implements the changes in the running-config file
2. Saves your changes to the startup-config file

If you decide not to save and implement the changes in the screen, select **Cancel** to discard them and continue switch operation with the current operation. For example, suppose you have made the changes shown below in the System Information screen:

To save and implement the changes for all parameters in this screen, press the **[Enter]** key, then press **[S]** (for **Save**). To cancel all changes, press the **[Enter]** key, then press **[C]** (for **Cancel**)

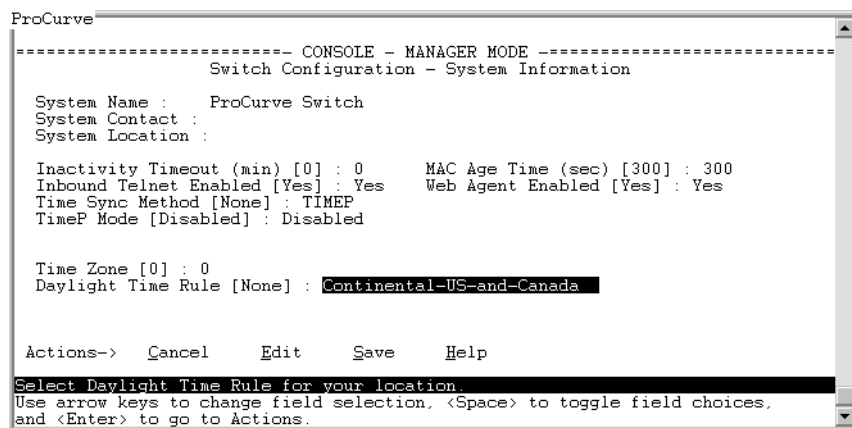


Figure 6-3. Example of Pending Configuration Changes You Can Save or Cancel

Note

If you reconfigure a parameter in the CLI and then go to the menu interface without executing a **write memory** command, those changes are stored only in the running configuration (even if you execute a Save operation in the menu interface). If you then execute a switch **boot** command in the menu interface, the switch discards the configuration changes made while using the CLI. To ensure that changes made while using the CLI are saved, execute **write memory** in the CLI before rebooting the switch.

Rebooting from the Menu Interface

- Terminates the current session and performs a reset of the operating system
- Activates any configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch. See “To Display the Port Counter Summary Report” on page B-12.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that the Reboot Switch option is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)

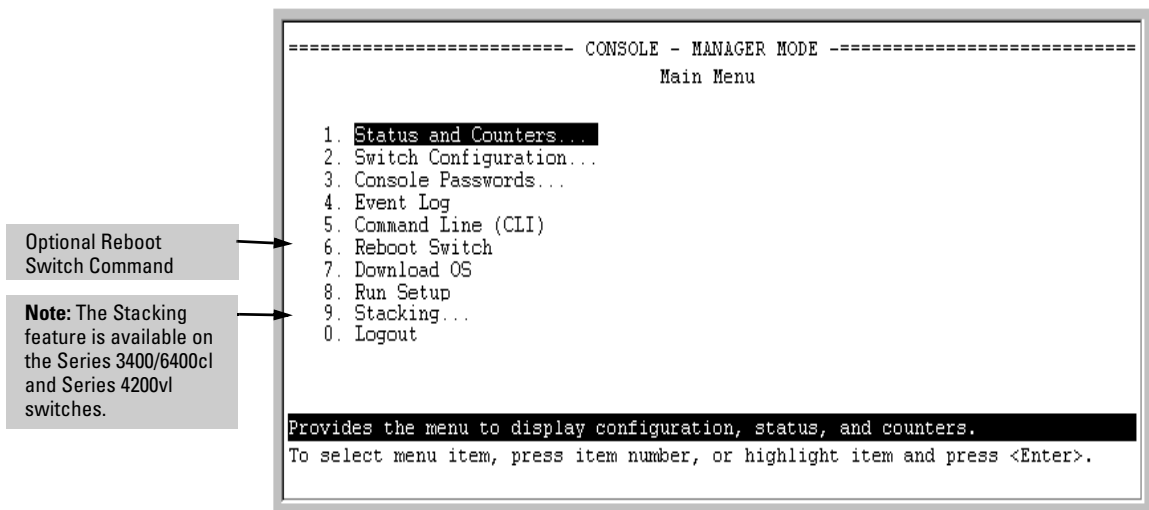


Figure 6-4. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support** parameter.

(To access these parameters, go to the Main menu and select **2. Switch Configuration**, then **8. VLAN Menu**, then **1. VLAN Support**.)

Switch Memory and Configuration

Using the Menu and Web Browser Interfaces To Implement Configuration Changes

If configuration changes requiring a reboot have been made, the switch displays an asterisk (*) next to the menu item in which the change has been made. For example, if you change and save parameter values for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** ..entry in the Main menu, as shown in figure 4-6:

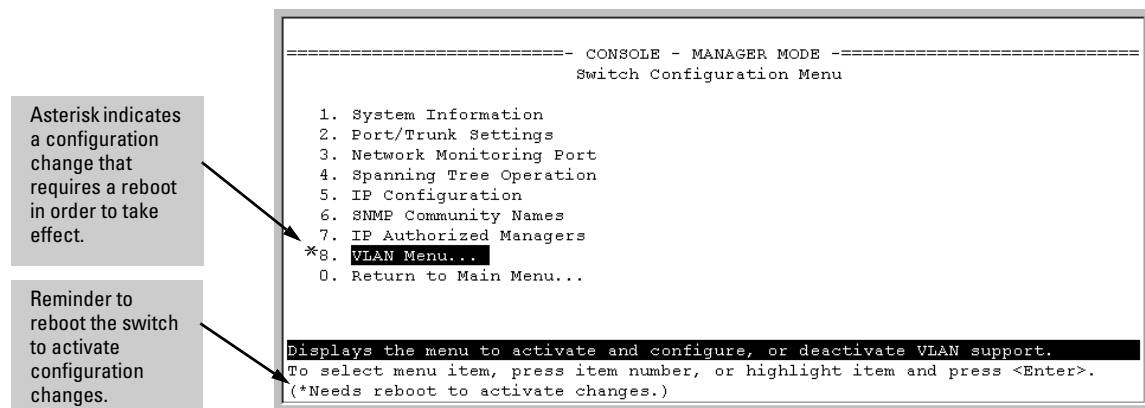


Figure 6-5. Indication of a Configuration Change Requiring a Reboot

Web: Implementing Configuration Changes

You can use the web browser interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change (in most cases, by clicking on **[Apply Changes]** or **[Apply Settings]**, you simultaneously change both the running-config file and the startup-config file.

Note

If you reconfigure a parameter in the CLI and then go to the browser interface without executing a **write memory** command, those changes will be saved to the startup-config file if you click on **[Apply Changes]** or **[Apply Settings]** in the web browser interface.

Using Primary and Secondary Flash Image Options

The Series switches covered by this guide feature two flash memory locations for storing switch software image files:

- **Primary Flash:** The default storage for a switch software image.
- **Secondary Flash:** The additional storage for either a redundant or an alternate switch software image.

With the Primary/Secondary flash option you can test a new image in your system without having to replace a previously existing image. You can also use the image options for troubleshooting. For example, you can copy a problem image into Secondary flash for later analysis and place another, proven image in Primary flash to run your system. The switch can use only one image at a time.

The following tasks involve primary/secondary flash options:

- Displaying the current flash image data and determining which switch software versions are available
- Switch software downloads
- Replacing and removing (erasing) a local switch software version
- System booting

Displaying the Current Flash Image Data

Use the commands in this section to:

- Determine whether there are flash images in both primary and secondary flash
- Determine whether the images in primary and secondary flash are the same
- Identify which switch software version is currently running

Viewing the Currently Active Flash Image Version. This command identifies the software version on which the switch is currently running, and whether the active version was booted from the primary or secondary flash image.

Syntax: show version

For example, if the switch is using a software version of E.08.22 stored in Primary flash, **show version** produces the following:

```
ProCurve(config)# show version
Image stamp:    /sw/code/build/info(s01)
                Dec 24 2005 10:50:26
                E.09.03
                1796
Boot Image:     Primary
```

Figure 6-6. Example Showing the Identity of the Current Flash Image (5300xl)

Determining Whether the Flash Images Are Different Versions. If the flash image sizes in primary and secondary are the same, then in almost every case, the primary and secondary images are identical. This command provides a comparison of flash image sizes, plus the boot ROM version and from which flash image the switch booted. For example, in the following case, the images are different versions of the switch software, and the switch is running on the version stored in the secondary flash image:

```
ProCurve(config)# show flash
Image          Size(Bytes)  Date    Version
-----
Primary Image  : 3275389    05/11/04 E.08.30
Secondary Image: 3258128    02/06/04 E.08.22
Boot Rom Version: E.05.04
Current Boot   : Primary
```

The unequal code size and differing dates indicate two different versions of the software.

Figure 6-7. Example Showing Different Flash Image Versions (5300xl)

Determining Which Flash Image Versions Are Installed. The **show version** command displays which software version the switch is currently running and whether that version booted from primary or secondary flash. Thus, if the switch booted from primary flash, you will see the version number of the software version stored in primary flash, and if the switch booted from secondary flash, you will see the version number of the software version stored in secondary flash. Thus, by using **show version**, then rebooting the switch from the opposite flash image and using **show version** again, you can determine the version(s) of switch software in both flash sources. For example:

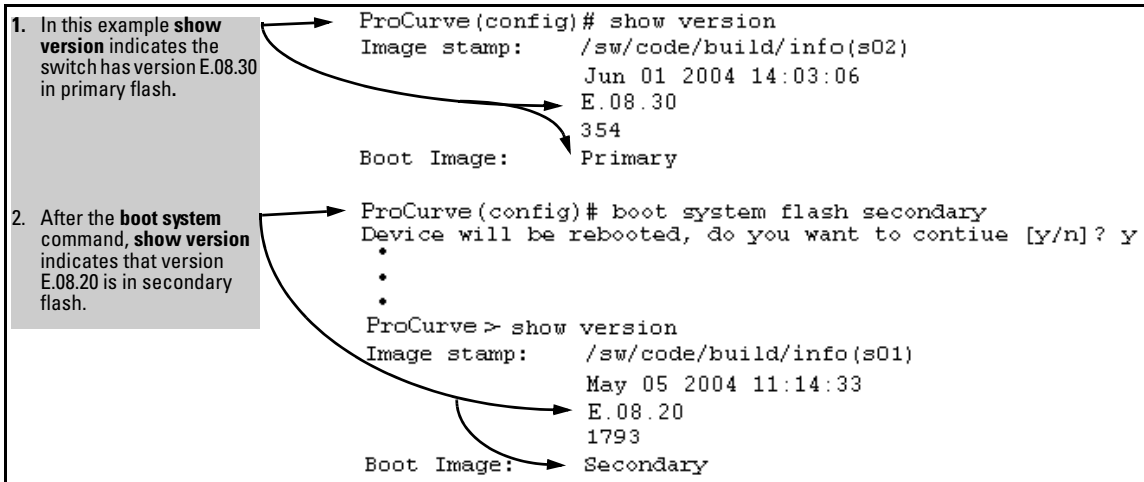


Figure 6-8. Determining the Software Version in Primary and Secondary Flash

Switch Software Downloads

The following table shows the switch's options for downloading a software version to flash and booting the switch from flash

Table 6-1. Primary/Secondary Memory Access

Action	Menu	CLI	Web Browser	SNMP
Download to Primary	Yes	Yes	Yes	Yes
Download to Secondary	No	Yes	No	Yes
Boot from Primary	Yes	Yes	Yes	Yes
Boot from Secondary	No	Yes	No	Yes

The different software download options involve different **copy** commands, plus **xmodem**, and **ftp**. These topics are covered in appendix A, "File Transfers".

Download Interruptions. In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted, as a result of an interruption, the switch will reboot from secondary flash and you can either copy the secondary image into primary or download another image to primary from an external source. See Appendix A, "File Transfers".

Local Switch Software Replacement and Removal

This section describes commands for erasing a software version and copying an existing software version between primary and secondary flash.

Note

It is not necessary to erase the content of a flash location before downloading another software file. The process automatically overwrites the previous file with the new file. If you want to remove an unwanted software version from flash, ProCurve recommends that you do so by overwriting it with the same software version that you are using to operate the switch, or with another acceptable software version. To copy a software file between the primary and secondary flash locations, see “Copying a Switch Software Image from One Flash Location to Another”, below.

The local commands described here are for flash image management within the switch. To download a software image file from an external source, refer to Appendix A, “File Transfers”.

Copying a Switch Software Image from One Flash Location to

Another. When you copy the flash image from primary to secondary or the reverse, the switch overwrites the file in the destination location with a copy of the file from the source location. This means you *do not* have to erase the current image at the destination location before copying in a new image.

Caution

Verify that there is an acceptable software version in the source flash location from which you are going to copy. Use the **show flash** command or, if necessary, the procedure under “Determining Which Flash Image Versions Are Installed” on page 6-14 to verify an acceptable software version. Attempting to copy from a source image location that has a corrupted flash image overwrites the image in the destination flash location. In this case, the switch will not have a valid flash image in either flash location, but will continue running on a temporary flash image in RAM. *Do not reboot the switch.* Instead, immediately download another valid flash image to primary or secondary flash. Otherwise, if the switch is rebooted without a software image in either primary or secondary flash, the temporary flash image in RAM will be cleared and the switch will go down. To recover, see “Restoring a Flash Image” on page C-58 (in the Troubleshooting chapter).

Syntax: copy flash flash <destination flash>

where: *destination flash* = **primary** or **secondary**:

For example, to copy the image in secondary flash to primary flash:

1. Verify that there is a valid flash image in the secondary flash location. The following figure indicates that a software image is present in secondary flash. (If you are unsure whether the image is secondary flash is valid, try booting from it before you proceed, by using **boot system flash secondary**.)

ProCurve(config)# show flash				The unequal code size, differing dates, and differing version numbers indicates two different versions of the software.
Image	Size(Bytes)	Date	Version	
-----	-----	-----	-----	
Primary Image	: 3275389	05/11/04	E.08.30	
Secondary Image	: 3258128	02/06/04	E.08.22	
Boot Rom Version:	E.05.04			
Current Boot	: Primary			

Figure 6-9. Example Indicating Two Different Software Versions in Primary and Secondary Flash

Execute the copy command as follows:

```
ProCurve(config)# copy flash flash primary
```

Erasing the Contents of Primary or Secondary Flash. This command deletes the software image file from the specified flash location.

Caution:

No Undo!

Before using this command in one flash image location (primary or secondary), ensure that you have a valid software file in the other flash image location (secondary or primary). If the switch has only one flash image loaded (in either primary or secondary flash) and you erase that image, then the switch does not have a software image stored in flash. In this case, if you do not reboot or power cycle the switch, you can recover by using xmodem or tftp to download another software image.

Syntax: erase flash < primary | secondary >

For example, to erase the software image in primary flash, do the following:

1. First verify that a usable flash image exists in secondary flash. The most reliable way to ensure this is to reboot the switch from the flash image you want to retain. For example, if you are planning to erase the primary image, then first reboot from the secondary image to verify that the secondary image is present and acceptable for your system:

```
ProCurve# boot system flash secondary
```

2. Then erase the software image in the selected flash (in this case, primary):

```
ProCurve# erase flash primary
The Primary OS Image will be deleted, continue [y/n]? _
```

The prompt shows which flash location will be erased.

Figure 6-10. Example of Erase Flash Prompt

3. Type **y** at the prompt to complete the flash erase.
4. Use **show flash** to verify erasure of the selected software flash image

```
ProCurve# show flash
Compressed Primary Code size    = 0
Compressed Secondary Code size  = 2555802
Boot Rom Version:               E.05.04
Current Boot:                   Secondary
```

The "0" here shows that primary flash has been erased.

Figure 6-11. Example of Show Flash Listing After Erasing Primary Flash

Rebooting the Switch

The switch offers reboot options through the **boot** and **reload** commands, plus the options inherent in a dual-flash image system. Generally, using **boot** provides more comprehensive self-testing; using **reload** gives you a faster reboot time.

Table 6-2. Comparing the Boot and Reload Commands

Actions	Included In Boot?	Included In Reload	Note
Save all configuration changes since the last boot or reload	Optional, with prompt	Yes, automatic	Config changes saved to the startup-config file
Perform all system self-tests	Yes	No	Reload provides a faster system reboot.
Choice of primary or secondary	Yes	No—Uses the current flash image.	

Booting from Primary Flash. This command always boots the switch from primary flash, executes the complete set of subsystem self-tests, and gives you the option of saving or discarding any configuration changes in the running-config file.

Syntax: boot

For example, to boot the switch from primary flash with pending configuration changes in the running-config file:

```
ProCurve(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
Boot from primary flash
Do you want to save current configuration [y/n]? _
```

Figure 6-12. Example of Boot Command (Default Primary Flash)

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation. (Entering **y** saves any configuration changes from the running-config file to the startup-config file; entering **n** discards them.) Also, if there are no pending configuration changes in the running-config file, then the reboot commences without the pause to display `Boot from primary flash`.

Booting from a Specified Flash. This version of the boot command gives you the option of specifying whether to reboot from primary or secondary flash, and is the required command for rebooting from secondary flash. This option also executes the complete set of subsystem self-tests.

Syntax: boot system flash < primary | secondary >

For example, to reboot the switch from secondary flash when there are no pending configuration changes in the running-config file:

```
ProCurve(config)# boot system flash secondary
Device will be rebooted, do you want to continue [y/n]? y
Boot from secondary flash
Do you want to save current configuration [y/n]? _
```

Figure 6-13. Example of Boot Command with Primary/Secondary Flash Option

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation. Also, if there are no pending configuration changes in the running-config file, then the reboot commences without the pause to display `Boot from secondary flash`.

Using the Fastboot feature. The **fastboot** command allows a boot sequence that skips the internal power-on self-tests, resulting in a faster boot time.

Syntax: [no] fastboot

Enables the fastboot option

[no]: *disables the feature.*

Syntax: show fastboot

Shows the status of the fastboot feature, either enabled or disabled.

The fastboot command is shown below.

```
ProCurve(config)# fastboot
```

Figure 6-14. Example of the Fastboot Command

Rebooting from the Current Software Version. **Reload** reboots the switch from the flash image and startup-config file on which the switch is currently running, and provides the option for saving to the startup-config file any configuration changes currently in the running-config file. Because **reload** bypasses some subsystem self-tests, the switch reboots faster than if you use either of the **boot** command options.

Syntax: reload

For example, if you change the number of VLANs the switch supports, you must reboot the switch in order to implement the change. Reload automatically saves your configuration changes and reboots the switch from the same software image you have been using:

```
ProCurve(config)# max-vlans 12
Command will take effect after saving configuration and reboot.
ProCurve(config)# reload
Device will be rebooted, do you want to continue [y/n]? y
Do you want to save current configuration [y/n]? _
```

Figure 6-15. Using Reload with Pending Configuration Changes

Operating Notes

Default Boot Source. The switch reboots from primary flash by default unless you specify the secondary flash.

Boot Attempts from an Empty Flash Location. In this case, the switch aborts the attempt and displays

```
Image does not exist  
Operation aborted.
```

Interaction of Primary and Secondary Flash Images with the Current Configuration. The switch has one startup-config file (page 6-3), which it always uses for reboots, regardless of whether the reboot is from primary or secondary flash. Also, for rebooting purposes, it is not necessary for the software image and the startup-config file to support identical software features. For example, suppose you have just downloaded a software upgrade that includes new features that are not supported in the software you used to create the current startup-config file. In this case, the software simply assigns factory-default values to the parameters controlling the new features. Similarly, If you create a startup-config file while using a version “Y” of the switch software, and then reboot the switch with an earlier software version “X” that does not include all of the features found in “Y”, the software simply ignores the parameters for any features that it does not support.

Multiple Configuration Files on 5300xl and 4200vl Switches

This section applies only to 5300xl switches running software release E.09.xx or greater, and 4200vl switches.

Action	Page
Listing and Displaying Startup-Config Files	6-27
Changing or Overriding the Reboot Configuration Policy	6-29
Managing Startup-Config Files	
Renaming Startup-Config Files	6-32
Copying Startup-Config Files	6-32
Erasing Startup-Config Files	6-34
Effect of Using the Clear + Reset Buttons	6-35
Copying Startup-Config Files to or from a Remote Server	6-37

With releases prior to E.09.xx, the 5300xl switch uses one startup-config file that automatically generates an identical configuration for the running-config file when the switch reboots. The switch operates using this running-config file, and any configuration changes affect only the running-config file until a **write memory** command is executed from the CLI (or a **Save** command is executed from the Menu interface) to copy the changes back to the startup-config file. Also, the same startup-config is used regardless of whether the switch reboots from the primary or secondary boot path.

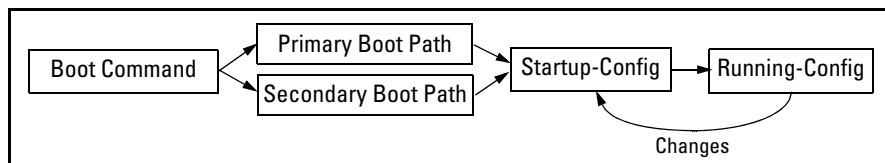


Figure 6-16. Reboot Process for Software Releases Earlier Than E.09.xx

This method of operation means that you cannot preserve different startup-config files across a reboot without using remote storage.

Beginning with software release E.09.xx, the switch allows up to three startup-config files with options for selecting which startup-config file to use for:

- A fixed reboot policy using a specific startup-config file for a specific boot path (primary or secondary flash)
- Overriding the current reboot policy on a per-instance basis

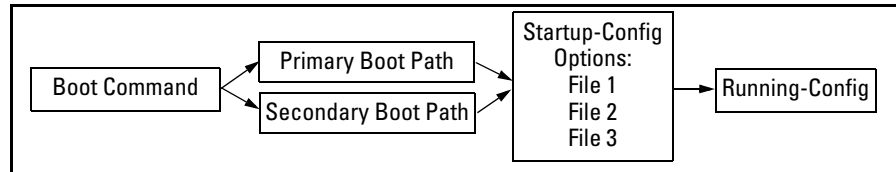


Figure 6-17. Optional Reboot Process for Software Release E.09.xx or Greater

While you can still use remote storage for startup-config files, you can now maintain multiple startup-config files on the switch and choose which version to use for a reboot policy or an individual reboot.

This choice of which configuration file to use for the startup-config at reboot provides the following new options:

- The switch can reboot with different configuration options without having to exchange one configuration file for another from a remote storage location.
- Transitions from one software release to another can be performed while maintaining a separate configuration for the different software release versions.
- By setting a reboot policy using a known good configuration and then overriding the policy on a per-instance basis, you can test a new configuration with the provision that if an unattended reboot occurs, the switch will come up with the known, good configuration instead of repeating a reboot with a misconfiguration.

General Operation

Multiple Configuration Storage in the Switch. The switch uses three memory “slots”, with identity (**id**) numbers of **1**, **2**, and **3**.

Memory Slots for Different Startup-Config Files	ProCurve(config)# show config files					
	Configuration files:					
	id	act	pri	sec	name	
	1				oldConfig	
	2	*	*	*	workingConfig	
	3					

A startup-config file stored in a memory slot has a unique, changeable file name. A software version earlier than release E.09.xx (that is, prior to the Multiple Configuration feature) always uses the startup-config file in memory slot 1, regardless of whether the software used for the reboot is stored in primary or secondary flash memory. Software version E.09.xx and greater can use the startup-config in any of the memory slots (if the software version supports the configured features).

Boot Options. With multiple startup-config files in the switch you can specify a policy for the switch to use upon reboot. The options include:

- Use the designated startup-config file with either or both reboot paths (primary or secondary flash)
- Override the current reboot policy for one reboot instance by specifying a boot path (primary or secondary flash) and the startup-config file to use.

Changing the Startup-Config File. When the switch reboots, the startup-config file supplies the configuration for the running-config file the switch uses to operate. Making changes to the running-config file and then executing a **write-mem** command (or, in the Menu interface, the **Save** command) are written back to the startup-config file used at the last reboot. For example, suppose that a system administrator performs the following on a switch that has two startup-config files (**workingConfig** and **backupConfig**):

1. Reboot the switch through the Primary boot path using the startup-config file named **backupConfig**.
2. Use the CLI to make configuration changes in the running-config file, and then execute **write mem**.

The result is that the startup-config file used to reboot the switch is modified by the actions in step 2.

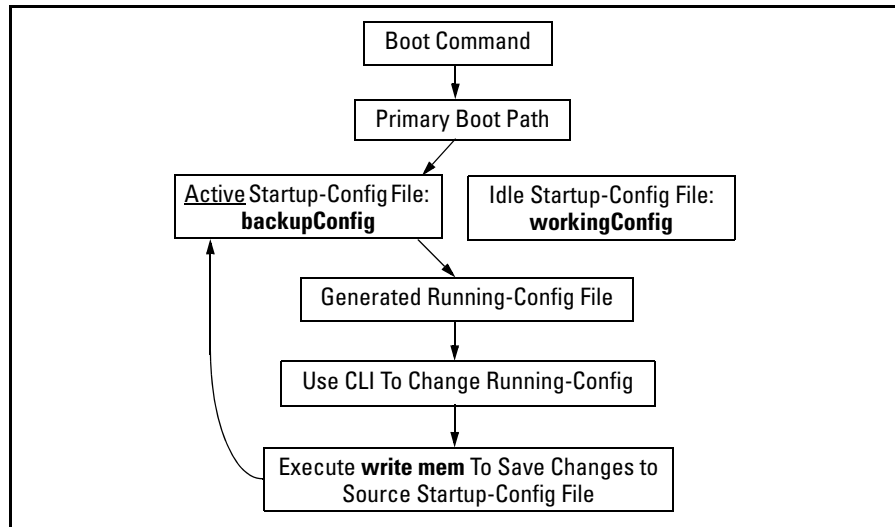


Figure 6-18. Example of Reboot Process and Making Changes to the Startup-Config File

Creating an Alternate Startup-Config File. There are two methods for creating a new configuration file:

- Copy an existing startup-config file to a new filename, then reboot the switch, make the desired changes to the running-config file, then execute **write memory**. (Refer to figure 6-18, above.)
- Erase the active startup-config file. This generates a new, default startup-config file that always results when the switch automatically reboots after deletion of the currently active startup-config file. (Refer to “Erasing a Startup-Config File” on page 6-34.)

Transitioning to Multiple Configuration Files

If your 5300xl switch was shipped from the factory with software release E.08.xx or earlier installed, you must download software release E.09.xx or greater to use the multiple configuration feature. At the first reboot with a software release supporting multiple configuration, the switch:

- Assigns the filename **oldConfig** to the existing startup-config file (which is stored in memory slot 1).
- Saves a copy of the existing startup-config file in memory slot 2 with the filename **workingConfig**.
- Assigns the **workingConfig** file as the active configuration and the default configuration for all subsequent reboots using either primary or secondary flash.

```
ProCurve(config)# show config files
```

Configuration files:

id	act	pri	sec	name
1				oldConfig
2	*	*	*	workingConfig
3				

Figure 6-19. Switch Memory Assignments After the First Reboot from Software Supporting Multiple Configuration

In the above state, the switch always:

- Uses the **workingConfig** file to reboot with E.09.xx or greater software releases
- Uses the **oldConfig** file to reboot with E.08.xx or earlier software

The commands described later in this section enable you to view the current multiple configuration status, manage multiple startup-config files, configure reboot policies, and override reboot policies on a per-instance basis.

Listing and Displaying Startup-Config Files

Command	Page
show config files	Below
show config < filename >	6-29

Viewing the Startup-Config File Status with Multiple Configuration Enabled

Rebooting the switch with software release E.09.xx or later automatically enables the multiple configuration feature.

Syntax: show config files

This command displays the available startup-config files on the switch and the current use of each file.

id: Identifies the memory slot for each startup-config file available on the switch. Software versions earlier than E.09.xx always use the startup-config file in slot 1.

act: An asterisk (*) in this column indicates that the corresponding startup-config file is currently in use.

pri: An asterisk (*) in this column indicates that the corresponding startup-config file is currently assigned to the primary boot path.

sec: An asterisk (*) in this column indicates that the corresponding startup-config file is currently assigned to the secondary boot path.

name: Shows the filename for each listed startup-config file in the switch. Refer to “Renaming an Existing Startup-Config File” on page 6-32 for the command you can use to change existing startup-config filenames.

— Continued on the next page. —

— Continued from the previous page. —

In the default configuration:

- *If the switch was shipped from the factory with software release E.09.xx installed in both the primary and secondary boot paths, then one startup-config file named **config1** is used for both paths and is stored in memory slot 1. Memory slots 2 and 3 are empty in this default configuration.)*
- *If the switch is running a software version earlier than E.09.xx and you download software version E.09.xx or greater to one of the boot paths, then the startup-config file in memory slot 1 supports the pre-E.09.xx software version and the startup-config file in memory slot 2 supports the E.09.xx (or greater) software version. (In this case, the default filename in memory slot 1 is **oldConfig** and the default filename in slot 2 is **workingConfig**. Memory slot 3 is empty in the default configuration.*

For example, after downloading software version E.09.xx or greater to the secondary flash in a 5300xl switch running software earlier than E.09.xx in primary flash, **show config files** displays the following:

ProCurve(config)# show config files				
Configuration files:				
id	act	pri	sec	name
1				oldConfig
2	*	*	*	workingConfig
3				

When you download software release E.09.xx for the first time and boot on this release, the switch places a copy of the **oldConfig** startup-config file in memory slot 2 and renames it **workingConfig**. This display shows that the **workingConfig** file in memory slot 2 and:

- Is the currently active configuration
- Is invoked when the switch boots with the software image in either the primary or secondary boot path.

As this example shows, you must reconfigure either the primary or the secondary boot path if you want to boot the switch using the startup-config file in another memory slot. (You can also change the above filenames. Refer to "Renaming an Existing Startup-Config File" on page 6-32.)

Figure 6-20. Example of Displaying the Current Multiple Configuration Status

Displaying the Content of A Specific Startup-Config File

With Multiple Configuration enabled, the switch can have up to three startup-config files. Because the **show config** command always displays the content of the currently active startup-config file, the command extension shown below is needed to allow viewing the contents of any other startup-config files stored in the switch.

Syntax: show config < filename >

*Available in software release E.09.xx or greater, this command displays the content of the specified startup-config file in the same way that the **show config** command displays the content of the default (currently active) startup-config file.*

Changing or Overriding the Reboot Configuration Policy

Command	Page
startup-default [primary secondary] config < filename >	Below
boot system flash < primary secondary > config < filename >	6-31

Prior to software release E.09.xx, the default boot configuration policy was to boot the switch using the current (only) startup-config file. Beginning with software release E.09.xx, which allows up to three separate startup-config files, you can boot the switch using any available startup-config file.

Changing the Reboot Configuration Policy. For a given reboot, the switch automatically reboots from the startup-config file assigned to the flash location (primary or secondary) being used for the current reboot. For example, when you first download a software version that supports multiple configuration files (E.09.xx or greater), and boot from the flash location of this version, the switch copies the existing startup-config file (named **oldConfig**) into memory slot 2, renames this file to **workingConfig**, and assigns **workingConfig** as:

- The active configuration file
- The configuration file to use when booting from either primary or secondary flash.

In this case, the switch is configured to automatically use the **workingConfig** file in memory slot 2 for all reboots. (Refer to figure 6-20 on page 6-28.)

You can use the following command to change the current policy so that the switch automatically boots using a different startup-config file.

Syntax: `startup-default [primary | secondary] config < filename >`

Specifies a boot configuration policy option:

[primary | secondary] config < filename >: *Designates the startup-config file to use in a reboot with the software version stored in a specific flash location. Use this option to change the reboot policy for either primary or secondary flash, or both.*

config < filename >: *Designates the startup-config file to use for all reboots, regardless of the flash version used. Use this option when you want to automatically use the same startup-config file for all reboots, regardless of the flash source used.*

Note: *To override the current reboot configuration policy for a single reboot instance, use the **boot system flash** command with the options described under “Overriding the Default Reboot Configuration Policy” on page 6-31.*

For example, suppose:

- Software release E.09.xx is stored in primary flash and a later software release is stored in secondary flash.
- The system operator is using memory slot 1 for a reliable, minimal configuration (named **minconfig**) for the software version in the primary flash, and slot 2 for a modified startup-config file (named **newconfig**) that includes untested changes for improved network operation with the software version in secondary flash.

The operator wants to ensure that in case of a need to reboot by pressing the Reset button, or if a power failure occurs, the switch will automatically reboot with the minimal startup-config file in memory slot 1. Since a reboot due to pressing the Reset button or to a power cycle always uses the software version in primary flash, the operator needs to configure the switch to always boot from primary flash with the startup-config file named **minconfig** (in memory slot 1). Also, whenever the switch boots from secondary flash, the operator also wants the startup-config named **newconfig** to be used. The following two commands configure the desired behavior.

```
ProCurve(config)# startup-default pri config minconfig
ProCurve(config) # startup-default sec config newconfig.
```

Overriding the Default Reboot Configuration Policy. This command provides a method for manually rebooting with a specific startup-config file other than the file specified in the default reboot configuration policy.

Syntax: boot system flash < primary | secondary > config < filename >

Specifies the name of the startup-config file to apply for the immediate boot instance only. This command overrides the current reboot policy. Software release E.09.xx or greater must be in the selected boot path (primary or secondary) to include config < filename > in the command string. (This command is not supported in earlier software releases.)

Using Reload To Reboot From the Current Flash Image and Startup-Config File.

Syntax: reload

*This command boots the switch from the currently active flash image and startup-config file. Because **reload** bypasses some subsystem self-tests, the switch boots faster than if you use a **boot** command.*

Note: To identify the currently active startup-config file, use the **show config files** command. For an example, refer to “Rebooting from the Current Software Version” on page 6-20.

Managing Startup-Config Files in the Switch

Command	Page
rename config < current-filename > < newname-str >	6-32
copy config < source-filename > config < dest-filename >	6-32
erase config < filename > startup-config	6-34
Erase startup-config using the front-panel Clear + Reset Buttons	6-35

Renaming an Existing Startup-Config File

Syntax: `rename config < current-filename > < newname-str >`

This command changes the name of an existing startup-config file. A file name can include up to 63, alphanumeric characters. Blanks are allowed in a file name enclosed in quotes (" " or ' '). (File names are not case-sensitive.)

Creating a New Startup-Config File

The switch allows up to three startup-config files. You can create a new startup-config file if there is an empty memory slot or if you want to replace one startup-config file with another.

Syntax: `copy config < source-filename > config < target-filename >`

This command makes a local copy of an existing startup-config file by copying the contents of an existing startup-config file in one memory slot to a new startup-config file in another, empty memory slot. This enables you to use a separate configuration file to experiment with configuration changes, while preserving the source file unchanged. It also simplifies a transition from one software version to another by enabling you to preserve the startup-config file for the earlier software version while creating a separate startup-config file for the later software version. With two such versions in place, you can easily reboot the switch with the correct startup-config file for either software version.

- *If the destination startup-config file already exists, it is overwritten by the content of the source startup-config file.*
- *If the destination startup-config file does not already exist, it will be created in the first empty configuration memory slot on the switch.*
- *If the destination startup-config file does not already exist, but there are no empty configuration memory slots on the switch, then a new startup-config file is not created and instead, the CLI displays the following error message:*

Unable to copy configuration to "< target-filename >".

For example, suppose both primary and secondary flash memory contain software release E.09.xx and use a startup-config file named **config1**:

```
ProCurve(config)# show config files
```

Configuration files:

id	act	pri	sec	name
1	*	*	*	config1
2				
3				

Figure 6-21. Example of Using One Startup-Config File for Both Primary and Secondary Flash

If you wanted to experiment with configuration changes to the software version in secondary flash, you could create and assign a separate startup-config file for this purpose.

```
ProCurve(config)# copy config config1 config config2
ProCurve(config)# startup-default secondary config config2
ProCurve(config)# show config files
```

Configuration files:

id	act	pri	sec	name
1	*	*		config1
2			*	config2
3				

The first two commands copy the **config1** startup-config file to **config2**, and then make **config2** the default startup-config file for booting from secondary flash.

Figure 6-22. Example of Creating and Assigning a New Startup-Config File

Note

You can also generate a new startup-config file by booting the switch from a flash memory location from which you have erased the currently assigned startup-config file. Refer to “Erasing a Startup-Config File” in the next section.

Erasing a Startup-Config File

You can erase any of the startup-config files in the switch's memory slots. In some cases, erasing a file causes the switch to generate a new, default-configuration file for the affected memory slot.

Syntax: `erase < config < filename >> | startup-config >`

config < filename >: *This option erases the specified startup-config file. If the specified file is not the currently active startup-config file, then the file is simply deleted from the memory slot it occupies. If the specified file is the currently active startup-config file, then the switch creates a new, default startup-config file with the same name as the erased file, and boots using this file. (This new startup-config file contains only the default configuration for the software version used in the reboot.)*

Note: *Where a file is assigned to either the primary or the secondary flash, but is not the currently active startup-config file, erasing the file does not remove the flash assignment from the memory slot for that file. Thus, if the switch boots using a flash location that does not have an assigned startup-config, then the switch creates a new, default startup-config file and uses this file in the reboot. (This new startup-config file contains only the default configuration for the software version used in the reboot.) Executing **write memory** after the reboot causes a switch-generated filename of **configx** to appear in the **show config files** display for the new file, where **x** corresponds to the memory slot number.*

startup-config: *This option erases the currently active startup-config file and reboots the switch from the currently active flash memory location. The erased startup-config file is replaced with a new startup-config file. The new file has the same filename as the erased file, but contains only the default configuration for the software version in the flash location (primary or secondary) used for the reboot. For example, suppose the last reboot was from primary flash using a configuration file named **minconfig**. Executing **erase startup-config** replaces the current content of **minconfig** with a default configuration and reboots the switch from primary flash.*

Figure 6-23 illustrates using **erase config <filename>** to remove a startup-config file.

```
ProCurve(config)# show config files
Configuration files:
  id | act pri sec | name
-----+-----+
  1  | *   *       | minconfig
  2  |           * | config2
  3  |           | config3

ProCurve(config)# erase config config3
ProCurve(config)# show config files
Configuration files:
  id | act pri sec | name
-----+-----+
  1  | *   *       | minconfig
  2  |           * | config2
  3  |           |
```

Figure 6-23. Example of Erasing a Non-Active Startup-Config File

With the same memory configuration as is shown in the bottom portion of figure 6-23, executing **erase startup-config** boots the switch from primary flash, resulting in a new file named **minconfig** in the same memory slot. The new file contains the default configuration for the software version currently in primary flash.

Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration

The Clear + Reset button combination described in the *Installation and Getting Started Guide* produces different results, depending on which software release is stored in primary flash. That is, when you press the Clear + Reset button combination:

- **Scenario 1:** If the primary flash location has a software release that does not support multiple configuration files (that is, earlier than E.09.xx) and the switch is running on software release E.09.xx or greater in secondary flash, then the switch:
 - Overwrites the content of the startup-config file currently in memory slot 1 with the default configuration for the software version in primary flash. (The filename does not change.)

- Boots the switch from primary flash using the new (default) configuration in the startup-config file in memory slot 1. Since the primary flash in this instance does not support multiple configuration files, the multiple configuration feature does not operate until the switch is booted again using software release E.09.xx or greater.

This scenario does not affect any startup-config files in memory slots 2 or 3, and does not change the active, primary, and secondary assignments the switch is maintaining in the configuration for software release E.09.xx or greater. (Note that the same applies if the switch is running on the software release earlier than E.09.xx in primary flash.)

- **Scenario 2:** If the primary flash location has a software version that supports multiple configuration files (release E.09.xx or greater), then the switch:
 - Overwrites the content of the startup-config file currently in memory slot 1 with the default configuration for the software version in primary flash, and renames this file to **config1**.
 - Erases any other startup-config files currently in memory.
 - Configures the new file in memory slot 1 as the default for both primary and secondary flash locations (regardless of the software version currently in secondary flash).
 - Boots the switch from primary flash using the new startup-config file.

ProCurve Switch 5304XL# sho config files				
Configuration files:				
id	act	pri	sec	name
1	*	*	*	config1
2				
3				

When the software version in the primary flash supports multiple configuration files, pressing Clear + Reset:

- Replaces all startup-config files with a single file named **config1** that contains the default configuration for the software version in primary flash.
- Resets the Active, Primary, and Secondary assignments as shown here.

Figure 6-24. Example of Clear + Reset Result from Scenario 2 on Page

Transferring Startup-Config Files To or From a Remote Server

Command	Page
copy config < src-file > tftp < ip-addr > < remote-file > < pc unix >	below
copy tftp config < dest-file > < ip-addr > < remote-file > < pc unix >	below
copy config < src-file > xmodem < pc unix >	6-38
copy xmodem config < dest-file > < pc unix >	6-38

TFTP: Copying a Configuration File to a Remote Host

Syntax: copy config < src-file > tftp < ip-addr > < remote-file > < pc | unix >

This is an addition to the copy tftp command options existing in releases prior to E.09.xx. Use this command to upload a configuration file from the switch to a TFTP server.

For more on using TFTP to copy a file to a remote server, refer to “TFTP: Copying a Configuration File to a Remote Host” on page A-24.

For example, the following command copies a startup-config file named **test-01** from the switch to a (UNIX) TFTP server at IP address 10.10.28.14:

```
ProCurve(config)# copy config test-01 tftp 10.10.28.14
test-01.txt unix
```

TFTP: Copying a Configuration File from a Remote Host

Syntax: copy tftp config < dest-file > < ip-addr > < remote-file > < pc | unix >

This is an addition to the copy tftp command options existing in releases prior to E.09.xx. Use this command to download a configuration file from a TFTP server to the switch.

Note: This command requires an empty memory slot in the switch. If there are no empty memory slots, the CLI displays the following message:

Unable to copy configuration to "< filename >".

For more on using TFTP to copy a file from a remote host, refer to “TFTP: Copying a Configuration from a Remote Host” on page A-23.

For example, the following command copies a startup-config file named **test-01.txt** from a (UNIX) TFTP server at IP address 10.10.28.14 to the first empty memory slot in the switch:

```
ProCurve(config)# copy tftp config test-01 10.10.28.14  
test-01.txt unix
```

Xmodem: Copying a Configuration File to a Serially Connected Host

Syntax: copy config < filename > xmodem < pc | unix >

*This is an addition to the **copy < config > xmodem** command options existing in releases prior to E.09.xx. Use this command to upload a configuration file from the switch to an Xmodem host.*

For more on using Xmodem to copy a file to a serially connected host, refer to “Xmodem: Copying a Configuration File from the Switch to a Serially Connected PC or UNIX Workstation” on page A-27.

Xmodem: Copying a Configuration from a Serially Connected Host

Syntax: copy xmodem config < dest-file > < pc | unix >

*This is an addition to the **copy xmodem** command options existing in releases prior to E.09.xx. Use this command to download a configuration file from an Xmodem host to the switch.*

For more on using Xmodem to copy a file from a serially connected host, refer to “Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation” on page A-27.

Operating Notes for Multiple Configuration Files

- SFTP/SCP: The configuration files are available for sftp/scp transfer as **/cfg/< filename >**.
- If you retain a software version earlier than E.09.*xxx* on the switch, always reserve the first config memory slot (**id = 1**) for a configuration compatible with the earlier version. This is because, software versions earlier than E.09.*xxx* always use the startup-config file in slot 1 to boot the switch.

Switch Memory and Configuration

Multiple Configuration Files on 5300xl and 4200vl Switches

—This page is intentionally unused—

Interface Access and System Information

Contents

Overview	7-2
Interface Access: Console/Serial Link, Web, and Inbound Telnet ..	7-3
Menu: Modifying the Interface Access	7-4
CLI: Modifying the Interface Access	7-5
Denying Interface Access by Terminating Remote Management Sessions	7-8
System Information	7-9
Menu: Viewing and Configuring System Information	7-10
CLI: Viewing and Configuring System Information	7-11
Web: Configuring System Parameters	7-14

Overview

This chapter describes how to:

- View and modify the configuration for switch interface access
- Use the CLI **kill** command to terminate a remote session
- View and modify switch system information

For help on how to actually use the interfaces built into the switch, refer to:

- Chapter 3, “Using the Menu Interface”
- Chapter 4, “Using the Command Line Interface (CLI)”
- Chapter 5, “Using the Web Browser Interface”

Why Configure Interface Access and System Information? The interface access features in the switch operate properly by default. However, you can modify or disable access features to suit your particular needs. Similarly, you can choose to leave the system information parameters at their default settings. However, modifying these parameters can help you to more easily distinguish one device from another in your network.

Interface Access: Console/Serial Link, Web, and Inbound Telnet

Interface Access Features

Feature	Default	Menu	CLI	Web
Inactivity Time	0 Minutes (disabled)	page 7-4	page 7-6	—
Inbound Telnet Access	Enabled	page 7-4	page 7-5	—
Outbound Telnet Access	n/a	—	page 7-6	—
Web Browser Interface Access	Enabled	page 7-4	page 7-6	—
Terminal type	VT-100	—	page 7-6	—
Event Log event types to list (Displayed Events)	All	—	page 7-6	—
Baud Rate	Speed Sense	—	page 7-6	—
Flow Control	XON/XOFF	—	page 7-6	—

In most cases, the default configuration is acceptable for standard operation.

Note

Basic switch security is through passwords. You can gain additional security by using the security features described in the Access Security Guide for your switch. You can also simply block unauthorized access via the web browser interface or Telnet (as described in this section) and installing the switch in a locked environment.

Menu: Modifying the Interface Access

The menu interface enables you to modify these parameters:

- Inactivity Timeout
- Inbound Telnet Enabled
- Web Agent Enabled

To Access the Interface Access Parameters:

1. From the Main Menu, Select...

2. Switch Configuration...

1. System Information

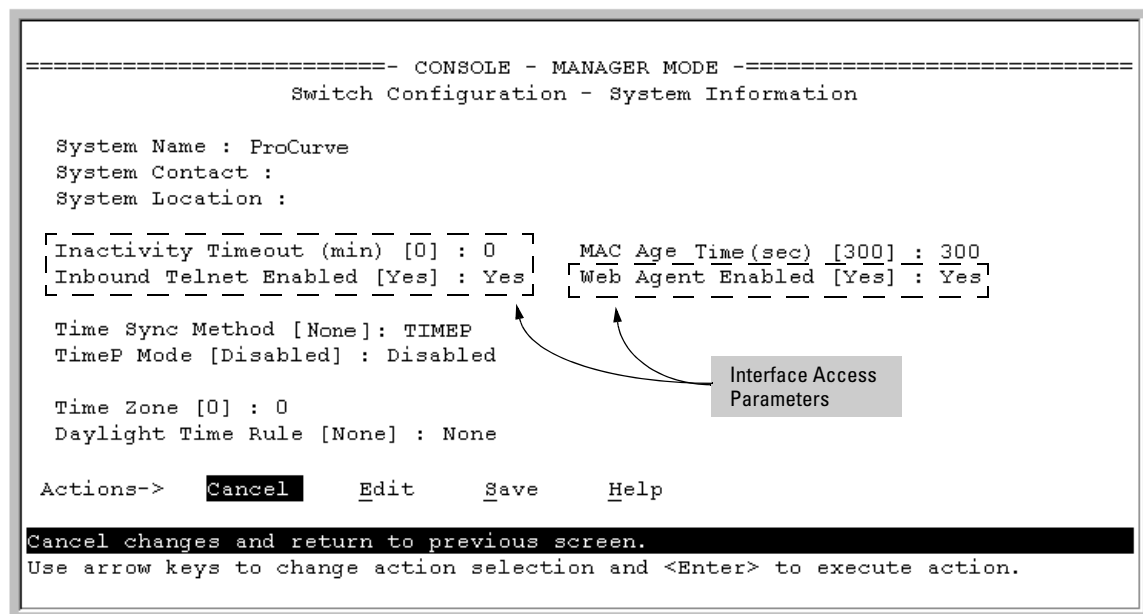


Figure 7-1. The Default Interface Access Parameters Available in the Menu Interface

2. Press [E] (for Edit). The cursor moves to the **System Name** field.
3. Use the arrow keys (J, U, L, R) to move to the parameters you want to change.

Refer to the online help provided with this screen for further information on configuration options for these features.

4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for Save).

CLI: Modifying the Interface Access

Interface Access Commands Used in This Section

show console	below
[no] telnet-server	below
[no] web-management	page 7-6
console	page 7-6

Listing the Current Console/Serial Link Configuration. This command lists the current interface access parameter settings.

Syntax: show console

This example shows the switch's default console/serial configuration.

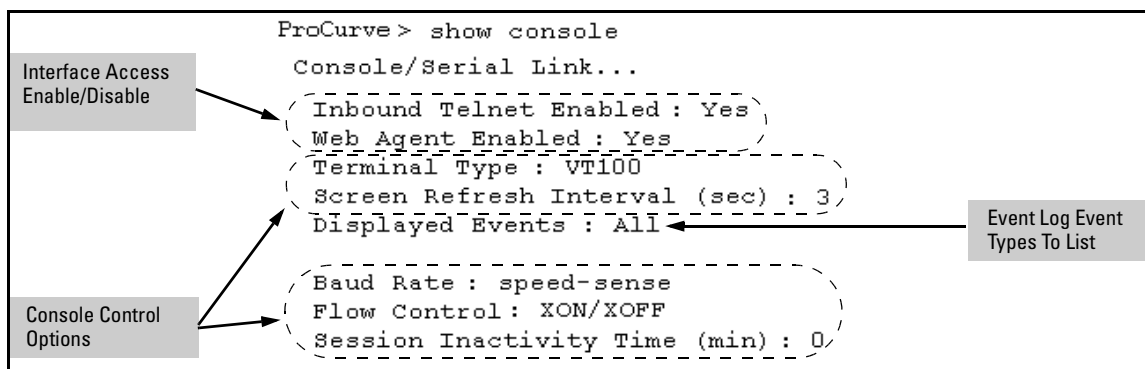


Figure 7-2. Listing of Show Console Command

Reconfigure Inbound Telnet Access. In the default configuration, inbound Telnet access is enabled.

Syntax: [no] telnet-server

To disable inbound Telnet access:

```
ProCurve(config)# no telnet-server
```

To re-enable inbound Telnet access:

```
ProCurve(config)# telnet-server
```

Outbound Telnet to Another Device. This feature operates independently of the telnet-server status and enables you to Telnet to another device that has an IP address.

Syntax: telnet < ip-address >

For example:

```
ProCurve # telnet 10.28.27.204
```

Reconfigure Web Browser Access. In the default configuration, web browser access is enabled.

Syntax: [no] web-management

To disable web browser access:

```
ProCurve(config)# no web-management
```

To re-enable web browser access:

```
ProCurve(config)# web-management
```

Reconfigure the Console/Serial Link Settings. You can reconfigure one or more console parameters with one console command.

Syntax: console

```
[terminal < vt100 | ansi | none >]
[screen-refresh < 1 | 3 | 5 | 10 | 20 | 30 | 45 | 60 >]
[baud-rate
  < speed-sense | 1200 | 2400 | 4800 | 9600 | 19200 | 38400 | 57600 |
  1155200 >]
[ flow-control < xon/xoff | none >]
[inactivity-timer < 0 | 1 | 5 | 10 | 15 | 20 | 30 | 60 | 120 >]
[events < none | all | non-info | critical | debug]
[local-terminal < vt 100 | none | ansi >]
```

Note

If you change the Baud Rate or Flow Control settings for the switch, you should make the corresponding changes in your console access device. Otherwise, you may lose connectivity between the switch and your terminal emulator due to differences between the terminal and switch settings for these two parameters.

All console parameter changes except **events** require that you save the configuration with **write memory** and then execute **boot** before the new console configuration will take effect.

For example, to use one command to configure the switch with the following:

- VT100 operation
- 19,200 baud
- No flow control
- 10-minute inactivity time
- Critical log events

you would use the following command sequence:

```
ProCurve(config)# console terminal vt100 baud-rate 19200 flow-control none
inactivity-timer 10 events critical
Command will take effect after saving configuration and reboot.
ProCurve(config)# write memory
ProCurve(config)# reload
```

The switch implements the EventLog change immediately. The switch implements the other console changes after executing **write memory** and **reload**.

Figure 7-3. Example of Executing the Console Command with Multiple Parameters

You can also execute a series of console commands and then save the configuration and boot the switch. For example:

Configure
the
individual
parameters.

```
ProCurve(config)# console baud-rate speed-sense
```

```
Command will take effect after saving configuration and reboot
```

```
ProCurve(config)# console flow-control xon/xoff
```

```
Command will take effect after saving configuration and reboot
```

```
ProCurve(config)# console inactivity-timer 0
```

```
Command will take effect after saving configuration and reboot
```

Save the
changes.

Boot the
switch.

```
ProCurve(config)# write memory
```

```
ProCurve(config)# reload
```

Figure 7-4. Example of Executing a Series of Console Commands

Denying Interface Access by Terminating Remote Management Sessions

The switch supports up to four management sessions. You can use **show ip ssh** to list the current management sessions, and **kill** to terminate a currently running remote session. (**Kill** does not terminate a Console session on the serial port, either through a direct connection or via a modem.)

Syntax: `kill [< session-number>]`

For example, if you are using the switch's serial port for a console session and want to terminate a currently active Telnet session, you would do the following:

```
ProCurve(config)# show ip ssh
  SSH Enabled           : Yes

  IP Port Number        : 22
  Timeout (sec)         : 120
  Server Key Size (bits) : 512

  Ses Type      Source IP and Port
  ---
  1  console
  2  telnet
  3  ssh      15.30.252.195:1531
  4  inactive

ProCurve(config)# kill 2
ProCurve(config)# show ip ssh
  SSH Enabled           : Yes

  IP Port Number        : 22
  Timeout (sec)         : 120
  Server Key Size (bits) : 512

  Ses Type      Source IP and Port
  ---
  1  console
  2  inactive
  3  ssh      15.30.252.195:1531
  4  inactive
```

Session 2 is an active Telnet session.

The kill 2 command terminates session 2.

Figure 7-5. Example of Using the “Kill” Command To Terminate a Remote Session

System Information

System Information Features

Feature	Default	Menu	CLI	Web
System Name	<i>switch product name</i>	page 7-10	page 7-12	page 7-14
System Contact	n/a	page 7-10	page 7-12	page 7-14
System Location	n/a	page 7-10	page 7-12	page 7-14
MAC Age Time	300 seconds	page 7-10	page 7-13	—
Time Sync Method	None	See Chapter 9, “Time Protocols”.		
Time Zone	0	page 7-10	page 7-13	—
Daylight Time Rule	None	page 7-10	page 7-13	—
Time	January 1, 1990 at 00:00:00 at last power reset	—	page 7-13	—

Configuring system information is optional, but recommended.

System Name: Using a unique name helps you to identify individual devices where you are using an SNMP network management tool such as ProCurve Manager.

System Contact and Location: This information is helpful for identifying the person administratively responsible for the switch and for identifying the locations of individual switches.

MAC Age Time: The number of seconds a MAC address the switch has learned remains in the switch’s address table before being aged out (deleted). Aging out occurs when there has been no traffic from the device belonging to that MAC address for the configured interval.

Time Sync Method: Selects the method (TimeP or SNTP) the switch will use for time synchronization. For more on this topic, refer to Chapter 9, “Time Protocols”.

Time Zone: The number of minutes your time zone location is to the West (+) or East (-) of Coordinated Universal Time (formerly GMT). The default **0** means no time zone is configured. For example, the time zone for Berlin, Germany is + 60 (minutes) and the time zone for Vancouver, Canada is - 480 (minutes).

Daylight Time Rule: Specifies the daylight savings time rule to apply for your location. The default is **None**. (For more on this topic, see appendix D, “Daylight Savings Time on ProCurve Switches.”)

Time: Used in the CLI to specify the time of day, the date, and other system parameters.

Menu: Viewing and Configuring System Information

To access the system information parameters:

1. From the Main Menu, Select...
2. **Switch Configuration...**
 1. **System Information**

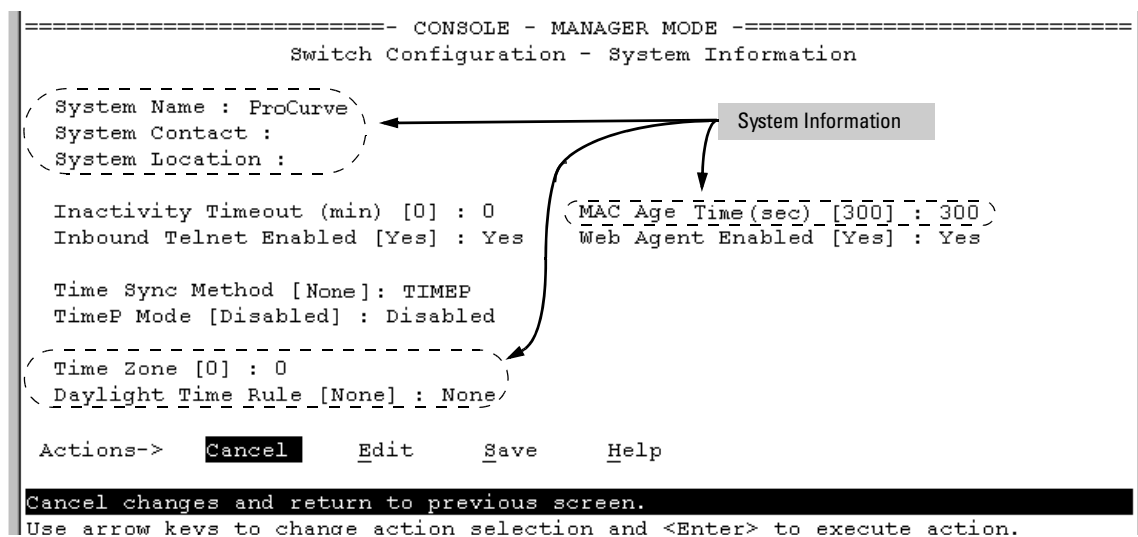


Figure 7-6. The System Information Configuration Screen (Default Values)

Note

To help simplify administration, it is recommended that you configure **System Name** to a character string that is meaningful within your system.

- 2. Press [E] (for Edit). The cursor moves to the **System Name** field.
- 3. Refer to the online help provided with this screen for further information on configuration options for these features.
- 4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for **S**ave) and return to the Main Menu.

CLI: Viewing and Configuring System Information

System Information Commands Used in This Section

show system-information	below
hostname	below
snmp-server [contact] [location]	below
mac-age-time	page 7-13
time	
timezone	page 7-13
daylight-time-rule	page 7-13
date	page 7-13
time	

Listing the Current System Information. This command lists the current system information settings.

Syntax: show system-information

This example shows the switch’s default console configuration.

```
ProCurve> show system-information
Status and Counters - General System Information
System Name       : HPswitch
System Contact    :
System Location   :
MAC Age Time(sec) : 300
Time Zone         : 0
Daylight Time Rule : None
```

Figure 7-7. Example of CLI System Information Listing

Configure a System Name, Contact, and Location for the Switch. To help distinguish one switch from another, configure a plain-language identity for the switch.

Syntax: `hostname <name-string>`
`snmp-server [contact <system-contact>] [location <system-location>]`

Both fields allow up to 48 characters. *Blank spaces* are not allowed in the variables for these commands.

For example, to name the switch “Blue” with “Ext-4474” as the system contact, and “North-Data-Room” as the location:

```
ProCurve(config)# hostname Blue
Blue(config)# snmp-server contact Ext-4474 location North-Data-Room
Blue(config)# show system-information

Status and Counters - General System Information
-----
System Name       : Blue
System Contact    : Ext-4474
System Location   : North-Data-Room
MAC Age Time (sec) : 300
Time Zone         : 0
Daylight Time Rule : None

Firmware revision : E.08.30      Base MAC Addr   : 0001e7-a0ec00
ROM Version       : E.05.04      Serial Number    : S000394041

Up Time           : 14 mins      Memory - Total   : 25,038,312
CPU Util (%)      : 1            Free             : 20,087,448

IP Mgmt - Pkts Rx : 0            Packet - Total   : 832
          Pkts Tx : 0            Buffers  Free     : 783
                                   Lowest    : 768

-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

New hostname, contact, and location data from previous commands.

Additional System Information

Figure 7-8. System Information Listing After Executing the Preceding Commands

Reconfigure the MAC Age Time for Learned MAC Addresses. This command corresponds to the MAC Age Interval in the menu interface, and is expressed in seconds.

Syntax: `mac-age-time < 10 - 1000000 > (seconds)`

For example, to configure the age time to seven minutes:

```
ProCurve(config)# mac-age-time 420
```

Configure the Time Zone and Daylight Time Rule. These commands:

- Set the time zone you want to use
- Define the daylight time rule for keeping the correct time when daylight-saving-time shifts occur.

Syntax: `time timezone < -720 - 840 >`
`time daylight-time-rule < none | alaska | continental-us-and-canada |`
`middle-europe-and-portugal | southern-hemisphere | western-europe |`
`user-defined>`

East of the 0 meridian, the sign is “+”. West of the 0 meridian, the sign is “-”.

For example, the time zone setting for Berlin, Germany is +60 (zone +1, or 60 minutes), and the time zone setting for Vancouver, Canada is -480 (zone -8, or -480 minutes). To configure the time zone and daylight time rule for Vancouver, Canada:

```
ProCurve(config)# time timezone -480  
daylight-time-rule continental-us-and-canada
```

Configure the Time and Date. The switch uses the time command to configure both the time of day and the date. Also, executing time without parameters lists the switch’s time of day and date. Note that the CLI uses a 24-hour clock scheme; that is, hour (*hh*) values from 1 p.m. to midnight are input as 13 - 24, respectively.

Syntax: `time [hh:mm [:ss]] [mm/dd/[yy] yy]`

For example, to set the switch to 9:45 a.m. on November 17, 2002:

```
ProCurve(config)# time 9:45 11/17/02
```

Note

Executing **reload** or **boot** resets the time and date to their default startup values.

Web: Configuring System Parameters

In the web browser interface, you can enter the following system information:

- System Name
- System Location
- System Contact

For access to the MAC Age Interval and the Time parameters, use the menu interface or the CLI.

Configure System Parameters in the Web Browser Interface.

1. Click on the **Configuration** tab.
2. Click on **[System Info]**.
3. Enter the data you want in the displayed fields.
4. Implement your new data by clicking on **[Apply Changes]**.

To access the web-based help provided for the switch, click on **[?]** in the web browser screen.

Configuring IP Addressing

Contents

Overview	8-2
IP Configuration	8-2
Just Want a Quick Start with IP Addressing?	8-3
IP Addressing with Multiple VLANs	8-4
Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL) ..	8-5
CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)	8-6
Web: Configuring IP Addressing	8-10
How IP Addressing Affects Switch Operation	8-11
DHCP/Bootp Operation	8-12
Network Preparations for Configuring DHCP/Bootp	8-14
IP Preserve: Retaining VLAN-1 IPAddressing	
Across Configuration File Downloads	8-15
Operating Rules for IP Preserve	8-16
Enabling IP Preserve	8-16

Overview

You can configure IP addressing through all of the switch's interfaces. You can also:

- Easily edit a switch configuration file to allow downloading the file to multiple switches without overwriting each switch's unique gateway and VLAN 1 IP addressing.
- Assign up to eight IP addresses to a VLAN (multinetting).

Why Configure IP Addressing? In its factory default configuration, the switch operates as a multiport learning bridge with network connectivity provided by the ports on the switch. However, to enable specific management access and control through your network, you will need IP addressing. Table 8-1 on page 8-11 shows the switch features that depend on IP addressing to operate.

IP Configuration

IP Configuration Features

Feature	Default	Menu	CLI	Web
IP Address and Subnet Mask	DHCP/Bootp	page 8-5	page 8-6	page 8-10
Multiple IP Addresses on a VLAN	n/a	—	page 8-8	—
Default Gateway Address	none	page 8-5	page 8-6	page 8-10
Packet Time-To-Live (TTL)	64 seconds	page 8-5	page 8-6	—
Time Server (Timep)	DHCP	page 8-5	page 8-6	—

IP Address and Subnet Mask. Configuring the switch with an IP address expands your ability to manage the switch and use its features. By default, the switch is configured to automatically receive IP addressing on the default VLAN from a DHCP/Bootp server that has been configured correctly with information to support the switch. (Refer to “DHCP/Bootp Operation” on page 8-12 for information on setting up automatic configuration from a server.) However, if you are not using a DHCP/Bootp server to configure IP addressing,

use the menu interface or the CLI to manually configure the initial IP values. After you have network access to a device, you can use the web browser interface to modify the initial IP configuration if needed.

For information on how IP addressing affects switch operation, refer to “How IP Addressing Affects Switch Operation” on page 8-11.

Multinetting: Assigning Multiple IP Addresses to a VLAN. For a given VLAN you can assign up to eight IP addresses. This allows you to combine two or more subnets on the same VLAN, which enables devices in the combined subnets to communicate normally through the network without needing to reconfigure the IP addressing in any of the combined subnets.

Default Gateway Operation. The default gateway is required when a router is needed for tasks such as reaching off-subnet destinations or forwarding traffic across multiple VLANs. The gateway value is the IP address of the next-hop gateway node for the switch, which is used if the requested destination address is not on a local subnet/VLAN. If the switch does not have a manually-configured default gateway and DHCP/Bootp is configured on the primary VLAN, then the default gateway value provided by the DHCP or Bootp server will be used. If the switch has a manually configured default gateway, then the switch uses this gateway, even if a different gateway is received via DHCP or Bootp on the primary VLAN. This is also true for manually configured TimeP, SNTP, and Time-To-Live(TTL). (In the default configuration, VLAN 1 is the Primary VLAN.) Refer to the information on Primary VLANs in the *Advanced Traffic Management Guide* for your switch.

Packet Time-To-Live (TTL) . This parameter specifies the maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet’s TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it. In most cases, the default setting (64) is adequate.

Just Want a Quick Start with IP Addressing?

If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using VLANs, ProCurve recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter **setup** at the CLI Manager level prompt.
`ProCurve# setup`
- Select **8. Run Setup** in the Main Menu of the menu interface.

For more on using the Switch Setup screen, see the *Installation and Getting Started Guide* you received with the switch.

IP Addressing with Multiple VLANs

In the factory-default configuration, the switch has one, permanent default VLAN (named DEFAULT_VLAN) that includes all ports on the switch. Thus, when only the default VLAN exists in the switch, if you assign an IP address and subnet mask to the switch, you are actually assigning the IP addressing to the DEFAULT_VLAN.

Notes

- If multiple VLANs are configured, then each VLAN can have its own IP address. This is because each VLAN operates as a separate broadcast domain and requires a unique IP address and subnet mask. A default gateway (IP) address for the switch is optional, but recommended.
- In the factory-default configuration, the default VLAN (named DEFAULT_VLAN) is the switch's *primary* VLAN. The switch uses the primary VLAN for learning the default gateway address. The switch can also learn other settings from a DHCP or Bootp server, such as (packet) Time-To-Live (TTL), and Timep or SNMP settings. (Other VLANs can also use DHCP or BootP to acquire IP addressing. However, the switch's gateway, TTL, and TimeP or SNTp values, which are applied globally, and not per-VLAN, will be acquired through the primary VLAN only, unless manually set by using the CLI, Menu, or web browser interface. (If these parameters are manually set, they will *not* be overwritten by alternate values received from a DHCP or Bootp server.) For more on VLANs, refer to the chapter titled "Static Virtual LANs" in the *Advanced Traffic Management Guide* for your switch.
- The IP addressing used in the switch should be compatible with your network. That is, the IP address must be unique and the subnet mask must be appropriate for your IP network.
- If you change the IP address through either Telnet access or the web browser interface, the connection to the switch will be lost. You can reconnect by either restarting Telnet with the new IP address or entering the new address as the URL in your web browser.

Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL)

Do one of the following:

- To manually enter an IP address, subnet mask, set the **IP Config** parameter to **Manual** and then manually enter the IP address and subnet mask values you want for the switch.
- To use DHCP or Bootp, use the menu interface to ensure that the **IP Config** parameter is set to **DHCP/Bootp**, then refer to “DHCP/Bootp Operation” on page 8-12.

To Configure IP Addressing.

1. From the Main Menu, Select.
 2. **Switch Configuration ...**
 5. **IP Configuration**

Notes

If multiple VLANs are configured, a screen showing all VLANs appears instead of the following screen.

The Menu interface displays the IP address for any VLAN. If you use the CLI to configure the IP address on a VLAN, use the CLI **show ip** command to list them. (Refer to “Viewing the Current IP Configuration” on page 8-6.)

For descriptions of these parameters, see the online Help for this screen.

Before using the DHCP/Bootp option, refer to “DHCP/Bootp Operation” on page 8-12.

```

=====-- CONSOLE - MANAGER MODE -----
                Switch Configuration - Internet (IP) Service

Default Gateway :
Default TTL      : 64

IP Config [DHCP/Bootp] : Manual
IP Address   : 15.30.248.184
Subnet Mask  : 255.255.248.0

Actions->  Cancel   Edit   Save   Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
  
```

Figure 8-1. Example of the IP Service Configuration Screen without Multiple VLANs Configured

2. Press [E] (for **Edit**).

3. If the switch needs to access a router, for example, to reach off-subnet destinations, select the **Default Gateway** field and enter the IP address of the gateway router.
4. If you need to change the packet Time-To-Live (TTL) setting, select **Default TTL** and type in a value between 2 and 255.
5. To configure IP addressing, select **IP Config** and do one of the following:
 - If you want to have the switch retrieve its IP configuration from a DHCP or Bootp server, at the **IP Config** field, keep the value as **DHCP/Bootp** and go to step 8.
 - If you want to manually configure the IP information, use the Space bar to select **Manual** and use the **[Tab]** key to move to the other IP configuration fields.
6. Select the **IP Address** field and enter the IP address for the switch.
7. Select the **Subnet Mask** field and enter the subnet mask for the IP address.
8. Press **[Enter]**, then **[S]** (for **Save**).

CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)

IP Commands Used in This Section	Page
show ip	8-6
ip address < mask-length >	8-7, 8-8
ip address /< mask-bits >	8-7, 8-8
ip default-gateway	8-10
ip ttl	8-10

Viewing the Current IP Configuration.

Syntax: show ip

This command displays the IP addressing for each VLAN configured in the switch. If only the DEFAULT_VLAN exists, then its IP configuration applies to all ports in the switch. Where multiple VLANs are configured, the IP addressing is listed per VLAN. The display includes switch-wide packet time-to-live, and (if configured) the switch's default gateway and Timestep configuration.

(You can also use the **show management** command to display the IP addressing and time server IP addressing configured on the switch. Refer to figure 9-6 on page 9-10.)

For example, in the factory-default configuration (no IP addressing assigned), the switch's IP addressing appears as:

The Default IP Configuration	ProCurve> show ip			
	Internet (IP) Service			
	Default Gateway :			
	Default TTL : 64			
	Arp Age : 20			
	TimeP Config : DHCP TimeP Poll Interval (min) : 720			
	VLAN	IP Config	IP Address	Subnet Mask
	-----	+	-----	-----
	DEFAULT_VLAN	DHCP/Bootp		

Figure 8-2. Example of the Switch's Default IP Addressing

With multiple VLANs and some other features configured, **show ip** provides additional information:

A Switch with IP Addressing and VLANs Configured	ProCurve> show ip			
	Internet (IP) Service			
	IP Routing : Disabled			
	Default Gateway : 10.28.227.1			
	Default TTL : 64			
	VLAN	IP Config	IP Address	Subnet Mask
	-----	+	-----	-----
	DEFAULT_VLAN	Manual	10.28.227.101	255.255.248.0
	VLAN_2	Disabled		

Figure 8-3. Example of Show IP Listing with Non-Default IP Addressing Configured

Configure an IP Address and Subnet Mask. The following command includes both the IP address and the subnet mask. You must either include the ID of the VLAN for which you are configuring IP addressing or go to the context configuration level for that VLAN. (If you are not using VLANs on the switch—that is, if the only VLAN is the default VLAN—then the VLAN ID is always “1”.)

Note

The default IP address setting for the DEFAULT_VLAN is **DHCP/Bootp**. On additional VLANs you create, the default IP address setting is **Disabled**.

Syntax: [no] vlan < vlan-id > ip address < ip-address/mask-length >
 or
 [no] vlan < vlan-id > ip address < ip-address > < mask-bits >
 or
 vlan < vlan-id > ip address dhcp-bootp

This example configures IP addressing on the default VLAN with the subnet mask specified in mask bits.

```
ProCurve(config)# vlan 1 ip address 10.28.227.103 255.255.255.0
```

This example configures the same IP addressing as the preceding example, but specifies the subnet mask by mask length.

```
ProCurve(config)# vlan 1 ip address 10.28.227.103/24
```

This example deletes an IP address configured in VLAN 1.

```
ProCurve (config) no vlan 1 ip address 10.28.227.103/24
```

Configure Multiple IP Addresses on a VLAN (Multinetting). You can configure up to eight IP addresses for the same VLAN. That is, the switch enables you to assign up to eight networks to a VLAN.

- Each IP address on a VLAN must be for a separate subnet.
- The switch allows up to 512 secondary subnet address assignments to VLANs.

Syntax: [no] vlan < vlan-id > ip address < ip-address/mask-length >
 [no] vlan < vlan-id > ip address < ip-address > < mask-bits >

For example, if you wanted to multinet VLAN_20 (VID = 20) with the IP addresses shown below, you would perform steps similar to the following. (For this example, assume that the first IP address is already configured.)

IP Address	VID	IP Address	Subnet Mask
1st address	20	10.25.33.101	255.255.240.0
2nd address	20	10.26.33.101	255.255.240.0
3rd address	20	10.27.33.101	255.255.240.0

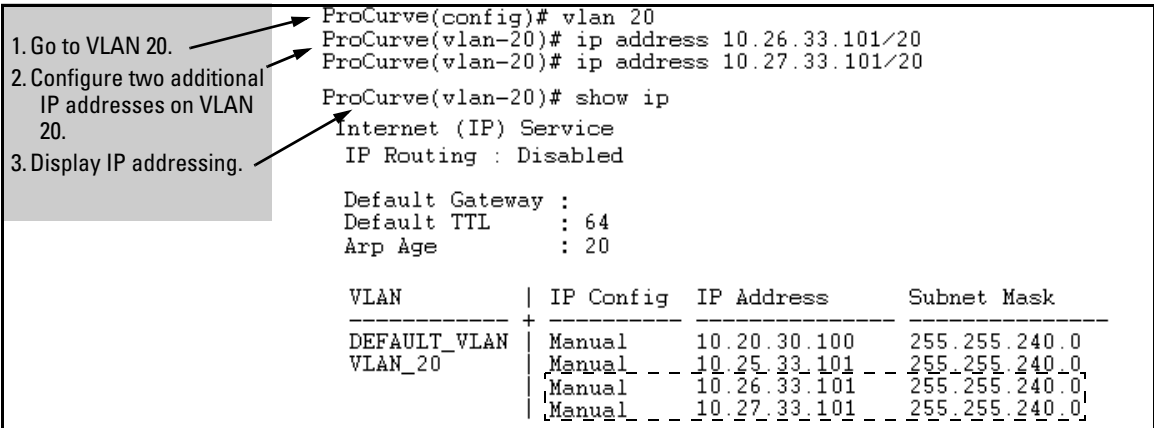


Figure 8-4. Example of Configuring and Displaying a Multinetted VLAN

If you then wanted to multinet the default VLAN, you would do the following:

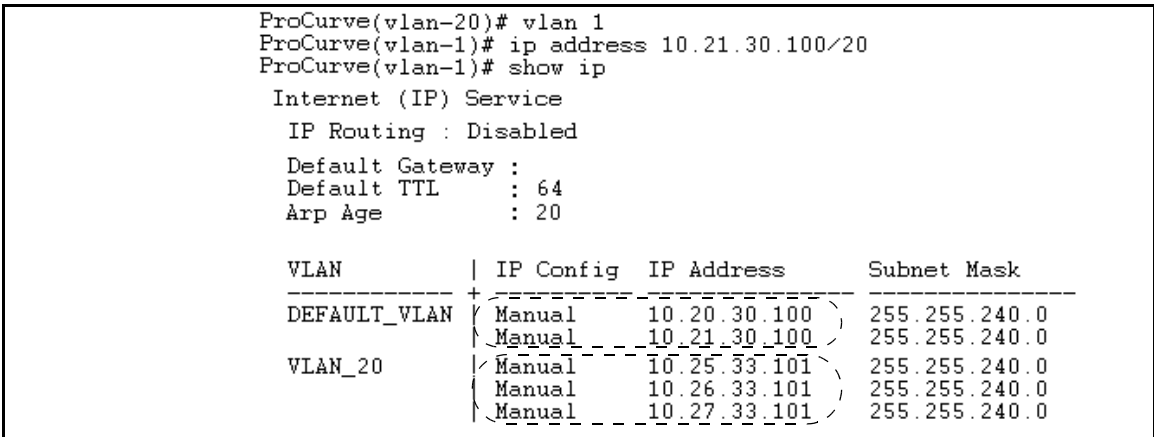


Figure 8-5. Example of Multinetting on the Default VLAN

Note

The Internet (IP) Service screen in the Menu interface (figure 8-1 on page 8-5) displays the first IP address for each VLAN. You must use the CLI **show ip** command to display the full IP address listing for multinetted VLANs.

Removing or Replacing IP Addresses in a Multinetted VLAN. To remove an IP address from a multinetted VLAN, use the **no** form of the IP address command shown on page 8-8. Generally, to replace one IP address with another, you should first remove the address you want to replace, and then enter the new address.

Configure the Optional Default Gateway. Using the Global configuration level, you can manually assign one default gateway to the switch. (The switch does *not* allow IP addressing received from a DHCP or Bootp server to replace a manually configured default gateway.)

Syntax: `ip default-gateway <ip-address>`

For example:

```
ProCurve(config)# ip default-gateway 10.28.227.115
```

Note

The switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. Thus, to avoid loss of Telnet access to off-subnet management stations, you should use the **ip route** command to configure a static (default) route before enabling routing. Refer to chapter 16, “IP Routing Features”, for more information.

Configure Time-To-Live (TTL). The maximum number of routers (hops) through which a packet can pass before being discarded. (The default is 64.) Each router decreases a packet’s TTL by 1 before forwarding the packet. If a router decreases the TTL to 0, the router drops the packet instead of forwarding it.

Syntax: `ip ttl <number-of-hops>`

```
ProCurve(config)# ip ttl 60
```

In the CLI, you can execute this command only from the global configuration level. The TTL default is 64, and the range is 2 - 255.

Web: Configuring IP Addressing

You can use the web browser interface to access IP addressing only if the switch already has an IP address that is reachable through your network.

1. Click on the Configuration tab.
2. Click on **[IP Configuration]**.

3. If you need further information on using the web browser interface, click on [?] to access the web-based help available for the switch.

How IP Addressing Affects Switch Operation

Without an IP address and subnet mask compatible with your network, the switch can be managed only through a direct terminal device connection to the Console RS-232 port. You can use direct-connect console access to take advantage of features that do not depend on IP addressing. However, to realize the full capabilities ProCurve proactive networking offers through the switch, configure the switch with an IP address and subnet mask compatible with your network. The following table lists the general features available with and without a network-compatible IP address configured.

Table 8-1. Features Available With and Without IP Addressing on the Switch

Features Available Without an IP Address	Additional Features Available with an IP Address and Subnet Mask
<ul style="list-style-type: none">• Direct-connect access to the CLI and the menu interface.• Stacking Candidate or Stack Member (Series 3400cl and Series 6400cl switches only)• DHCP or Bootp support for automatic IP address configuration, and DHCP support for automatic Timep server IP address configuration• Spanning Tree Protocol• Port settings and port trunking• Switch meshing• Console-based status and counters information for monitoring switch operation and diagnosing problems through the CLI or menu interface.• VLANs and GVRP• Serial downloads of software updates and configuration files (Xmodem)• Link test• Port monitoring• Password authentication• Quality of Service (QoS)• Authorized IP manager security	<ul style="list-style-type: none">• Web browser interface access, with configuration, security, and diagnostic tools, plus the Alert Log for discovering problems detected in the switch along with suggested solutions• SNMP network management access such as ProCurve Manager for network configuration, monitoring, problem-finding and reporting, analysis, and recommendations for changes to increase control and uptime• TACACS+, RADIUS, SSH, SSL, and 802.1x authentication• Multinetting on VLANs• Stacking Commander*• Telnet access to the CLI or the menu interface• IGMP• TimeP and SNTP server configuration• TFTP download of configurations and software updates• Access Control Lists (ACLs)• IP routing, Multicast Routing• XRRP router redundancy• PIM-DM (Series 5300xl switches only)• NAT (Series 5300xl switches only)• Ping test

*Although a Commander can operate without an Ip address, doing so makes it unavailable for in-band access in an IP network.

DHCP/Bootp Operation

Overview. DHCP/Bootp is used to provide configuration data from a DHCP or Bootp server to the switch. This data can be the IP address, subnet mask, default gateway, Timep Server address, and TFTP server address. If a TFTP server address is provided, this allows the switch to TFTP a previously saved configuration file from the TFTP server to the switch. With either DHCP or Bootp, the servers must be configured prior to the switch being connected to the network.

Note

The switches covered by this guide are compatible with both DHCP and Bootp servers.

The DHCP/Bootp Process. Whenever the **IP Config** parameter in the switch or in an individual VLAN in the switch is configured to **DHCP/Bootp** (the default), or when the switch is rebooted with this configuration:

1. DHCP/Bootp requests are automatically broadcast on the local network. (The switch sends one type of request to which either a DHCP or Bootp server can respond.)
2. When a DHCP or Bootp server receives the request, it replies with a previously configured IP address and subnet mask for the switch. The switch also receives an IP Gateway address if the server has been configured to provide one. In the case of Bootp, the server must first be configured with an entry that has the switch's MAC address. (To determine the switch's MAC address, see appendix D, "MAC Address Management".) The switch properly handles replies from either type of server. If multiple replies are returned, the switch tries to use the first reply.)

Note

If you manually configure default gateway, TTL, TimeP, and/or SNTP parameters on the switch, it ignores any values received for the same parameters via DHCP or Bootp.

If the switch is initially configured for DHCP/Bootp operation (the default), or if it reboots with this configuration, it begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process immediately.

DHCP Operation. A significant difference between a DHCP configuration and a Bootp configuration is that an IP address assignment from a DHCP server is automatic. Depending on how the DHCP server is configured, the switch may receive an IP address that is temporarily leased. Periodically the switch may be required to renew its lease of the IP configuration. Thus, the IP addressing provided by the server may be different each time the switch reboots or renews its configuration from the server. However, you can fix the address assignment for the switch by doing either of the following:

- Configure the server to issue an “infinite” lease.
- Using the switch’s MAC address as an identifier, configure the server with a “Reservation” so that it will always assign the same IP address to the switch. (For MAC address information, refer to appendix D, “MAC Address Management”.)

For more information on either of these procedures, refer to the documentation provided with the DHCP server.

Bootp Operation. When a Bootp server receives a request it searches its Bootp database for a record entry that matches the MAC address in the Bootp request from the switch. If a match is found, the configuration data in the associated database record is returned to the switch. For many Unix systems, the Bootp database is contained in the **/etc/bootptab** file. In contrast to DHCP operation, Bootp configurations are always the same for a specific receiving device. That is, the Bootp server replies to a request with a configuration previously stored in the server and designated for the requesting device.

Bootp Database Record Entries. A minimal entry in the Bootp table file **/etc/bootptab** to update an IP address and subnet mask to the switch or a VLAN configured in the switch would be similar to this entry:

```
5300switch:\
    ht=ether:\
    ha=0030c1123456:\
    ip=10.66.77.88:\
    sm=255.255.248.0:\
    gw=10.66.77.1:\
    hn:\
    vm=rfc1048
```

An entry in the Bootp table file `/etc/bootptab` to tell the switch or VLAN where to obtain a configuration file download would be similar to this entry:

```
5300switch:\
  ht=ether:\
  ha=0030c1123456:\
  ip=10.66.77.88:\
  sm=255.255.248.0:\
  gw=10.66.77.1:\
  lg=10.22.33.44:\
  T144="switch.cfg":\
  vm=rfc1048
```

where:

5300switch	is a user-defined symbolic name to help you find the correct section of the bootptab file. If you have multiple switches that will be using Bootp to get their IP configuration, you should use a unique symbolic name for each switch.
ht	is the "hardware type". For the switches covered in this guide, enter ether (for Ethernet). <i>This tag must precede the ha tag.</i>
ha	is the "hardware address". Use the switch's (or VLAN's) 12-digit MAC address.
ip	is the IP address to be assigned to the switch (or VLAN).
sm	is the subnet mask of the subnet in which the switch (or VLAN) is installed.
gw	is the IP address of the default gateway.
lg	TFTP server address (source of final configuration file)
T144	is the vendor-specific "tag" identifying the configuration file to download.
vm	is a required entry that specifies the Bootp report format. Use rfc1048 for the switches covered in this guide.

Note

The above Bootp table entry is a sample that will work for the switch when the appropriate addresses and file names are used.

Network Preparations for Configuring DHCP/Bootp

In its default configuration, the switch is configured for DHCP/Bootp operation. However, the DHCP/Bootp feature will not acquire IP addressing for the switch unless the following tasks have already been completed:

- For Bootp operation:
 - A Bootp database record has already been entered into an appropriate Bootp server.
 - The necessary network connections are in place
 - The Bootp server is accessible from the switch

- For DHCP operation:
 - A DHCP scope has been configured on the appropriate DHCP server.
 - The necessary network connections are in place
 - A DHCP server is accessible from the switch

Note

Designating a primary VLAN other than the default VLAN affects the switch's use of information received via DHCP/Bootp. For more on this topic, refer to the chapter describing VLANs in the *Advanced Traffic Management Guide* for your switch.

After you reconfigure or reboot the switch with DHCP/Bootp enabled in a network providing DHCP/Bootp service, the switch does the following:

- Receives an IP address and subnet mask and, if configured in the server, a gateway IP address and the address of a Tftp server.
- If the DHCP/Bootp reply provides information for downloading a configuration file, the switch uses TFTP to download the file from the designated source, then reboots itself. (This assumes that the switch or VLAN has connectivity to the TFTP file server specified in the reply, that the configuration file is correctly named, and that the configuration file exists in the TFTP directory.)

IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads

For the switches covered in this guide, IP Preserve enables you to copy a configuration file to multiple switches while retaining the individual IP address and subnet mask on VLAN 1 in each switch, and the Gateway IP address assigned to the switch. This enables you to distribute the same configuration file to multiple switches without overwriting their individual IP addresses.

Operating Rules for IP Preserve

When **ip preserve** is entered as the last line in a configuration file stored on a TFTP server:

- If the switch's current IP address for VLAN 1 was not configured by DHCP/Bootp, IP Preserve retains the switch's current IP address, subnet mask, and IP gateway address when the switch downloads the file and reboots. The switch adopts all other configuration parameters in the configuration file into the startup-config file.
- If the switch's current IP addressing for VLAN 1 is from a DHCP server, IP Preserve is suspended. In this case, whatever IP addressing the configuration file specifies is implemented when the switch downloads the file and reboots. If the file includes DHCP/Bootp as the IP addressing source for VLAN 1, the switch will configure itself accordingly and use DHCP/Bootp. If instead, the file includes a dedicated IP address and subnet mask for VLAN 1 and a specific gateway IP address, then the switch will implement these settings in the startup-config file.
- The **ip preserve** statement does not appear in **show config** listings. To verify IP Preserve in a configuration file, open the file in a text editor and view the last line. For an example of implementing IP Preserve in a configuration file, see figure 8-6, below.

Enabling IP Preserve

To set up IP Preserve, enter the **ip preserve** statement at the end of a configuration file. (Note that you do not execute IP Preserve by entering a command from the CLI).

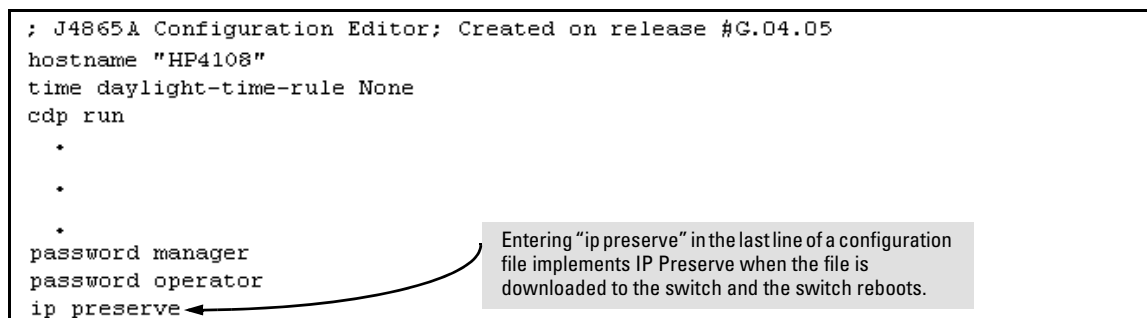


Figure 8-6. Example of Implementing IP Preserve in a Configuration File

For example, consider Figure 8-7:

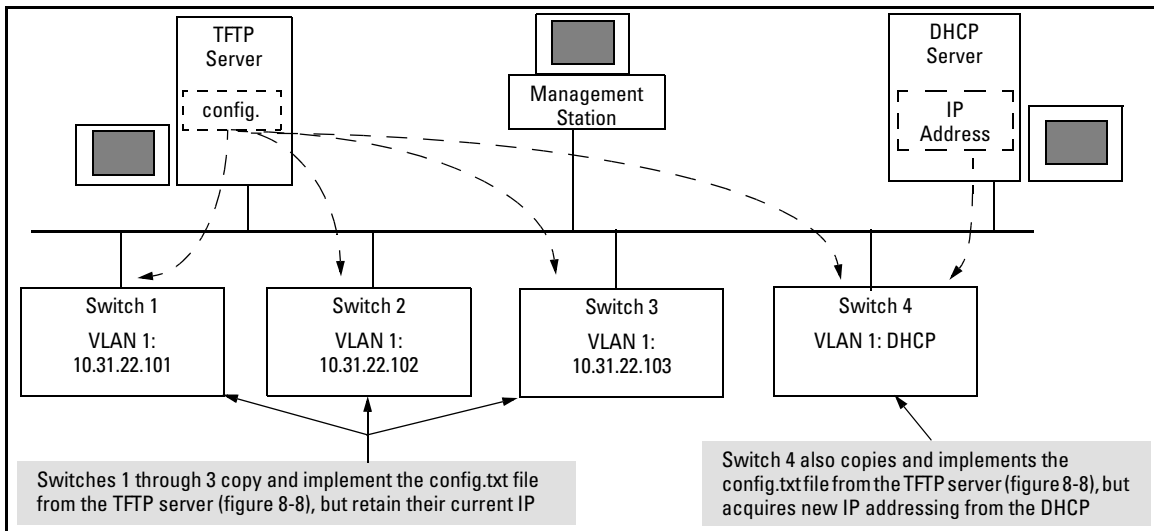


Figure 8-7. Example of IP Preserve Operation with Multiple Series Switches

If you apply the following configuration file to figure 8-7, switches 1 - 3 will retain their manually assigned IP addressing and switch 4 will be configured to acquire its IP addressing from a DHCP server.

```
; J4850A Configuration Editor; Created on release #E.08.22
hostname "HP4108"
time daylight-time-rule None
cdp run
interface A11
  no lACP
exit
interface A12
  no lACP
exit
trunk A11-A12 Trk1 Trunk
ip default-gateway 10.33.32.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  ip address dhcp-bootp
  exit
password manager
password operator
ip preserve
```

Using figure 8-7, above, switches 1 - 3 ignore these entries because the file implements IP Preserve and their current IP addressing was not acquired through DHCP/Bootp.

Switch 4 ignores IP Preserve and implements the DHCP/Bootp addressing and IP Gateway specified in this file (because its last IP addressing was acquired from a DHCP/Bootp server).

IP Preserve Command

Figure 8-8. Configuration File in TFTP Server, with DHCP/Bootp Specified as the IP Addressing Source

Configuring IP Addressing

IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads

If you apply this configuration file to figure 8-7, switches 1 - 3 will still retain their manually assigned IP addressing. However, switch 4 will be configured with the IP addressing included in the file.

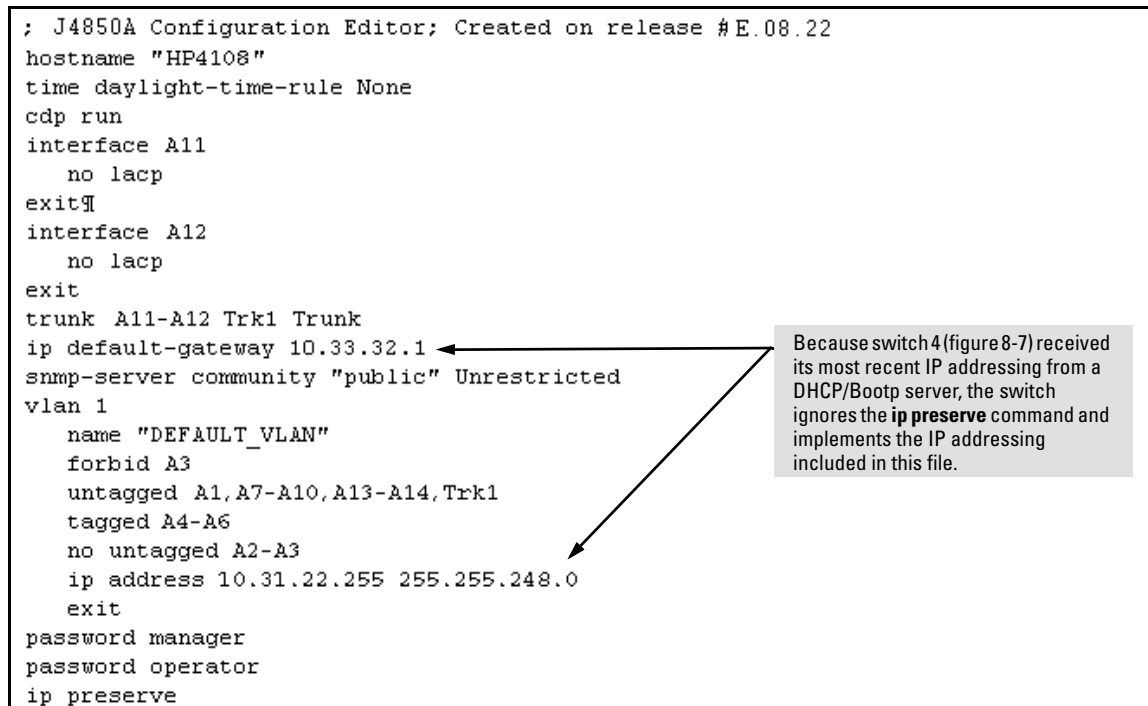


Figure 8-9. Configuration File in TFTP Server, with Dedicated IP Addressing Instead of DHCP/Bootp

To summarize the IP Preserve effect on IP addressing:

- If the switch received its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it ignores the IP Preserve command when it downloads the configuration file, and implements whatever IP addressing instructions are in the configuration file.
- If the switch did not receive its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it retains its current IP addressing when it downloads the configuration file.
- The content of the downloaded configuration file determines the IP addresses and subnet masks for other VLANs.

Time Protocols

Contents

Overview	9-2
TimeP Time Synchronization	9-2
SNTP Time Synchronization	9-3
Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation	9-3
General Steps for Running a Time Protocol on the Switch:	9-3
Disabling Time Synchronization	9-4
SNTP: Viewing, Selecting, and Configuring	9-5
Menu: Viewing and Configuring SNTP	9-6
CLI: Viewing and Configuring SNTP	9-8
Viewing the Current SNTP Configuration	9-9
Configuring (Enabling or Disabling) the SNTP Mode	9-10
TimeP: Viewing, Selecting, and Configuring	9-16
Menu: Viewing and Configuring TimeP	9-17
CLI: Viewing and Configuring TimeP	9-18
Viewing the Current TimeP Configuration	9-19
Configuring (Enabling or Disabling) the TimeP Mode	9-20
SNTP Unicast Time Polling with Multiple SNTP Servers	9-25
Address Prioritization	9-25
Displaying All SNTP Server Addresses Configured on the Switch ..	9-26
Adding and Deleting SNTP Server Addresses	9-26
Menu: Operation with Multiple SNTP Server Addresses Configured	9-28
SNTP Messages in the Event Log	9-28

Overview

This chapter describes:

- SNTP Time Protocol Operation
- Timep Time Protocol Operation

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

The switch offers TimeP and SNTP (Simple Network Time Protocol) and a **timesync** command for changing the time protocol selection (or turning off time protocol operation).

Notes

- Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time.
 - In the factory-default configuration, the time synchronization option is set to TimeP, with the TimeP mode itself set to **Disabled**.
-

TimeP Time Synchronization

You can either manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one, designated Timep server. This option enhances security by specifying which time server to use.

SNTP Time Synchronization

SNTP provides two operating modes:

- **Broadcast Mode:** The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address. Refer to the documentation provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable **Poll Interval** expires three consecutive times without an update received from the first-detected server.

Note

To use Broadcast mode, the switch and the SNTP server must be in the same subnet.

- **Unicast Mode:** The switch requests a time update from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI **sntp server** command.) This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.

Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

General Steps for Running a Time Protocol on the Switch:

1. Select the time synchronization protocol: **SNTP** or **TimeP** (the default).
2. Enable the protocol. The choices are:
 - SNTP: **Broadcast** or **Unicast**
 - TimeP: **DHCP** or **Manual**

Time Protocols

Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

3. Configure the remaining parameters for the time protocol you selected.

The switch retains the parameter settings for both time protocols even if you change from one protocol to the other. Thus, if you select a time protocol, the switch uses the parameters you last configured for the selected protocol.

Note that simply selecting a time synchronization protocol does not enable that protocol on the switch unless you also enable the protocol itself (step 2, above). For example, in the factory-default configuration, TimeP is the selected time synchronization method. However, because TimeP is disabled in the factory-default configuration, no time synchronization protocol is running.

Disabling Time Synchronization

You can use either of the following methods to disable time synchronization without changing the Timep or SNTP configuration:

- In the System Information screen of the Menu interface, set the **Time Synch Method** parameter to **None**, then press **[Enter]**, then **[S]** (for **Save**).
- In the Global config level of the CLI, execute **no timesync**.

SNTP: Viewing, Selecting, and Configuring

SNTP Feature	Default	Menu	CLI	Web
view the SNTP time synchronization configuration	n/a	page 9-6	page 9-9	—
select SNTP as the time synchronization method	timep	page 9-6	page 9-10 ff.	—
disable time synchronization	timep	page 9-6	page 9-14	—
enable the SNTP mode (Broadcast, Unicast, or Disabled)	disabled			—
broadcast	n/a	page 9-7	page 9-11	—
unicast	n/a	page 9-7	page 9-11	—
none/disabled	n/a	page 9-7	page 9-14	—
configure an SNTP server address (for Unicast mode only)	none	page 9-7	page 9-11 ff.	—
change the SNTP server version (for Unicast mode only)	3	page 9-7	page 9-13	—
change the SNTP poll interval	720 seconds	page 9-7	page 9-13	—

Table 9-1. SNTP Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either SNTP, TIMEP, or None as the time synchronization method.
SNTP Mode	
Disabled	The Default. SNTP does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI timesync command.
Unicast	Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address.
Broadcast	Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, it the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.

SNTP Parameter	Operation
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI. See “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 25.
Server Version	Default: 3; range: 1 - 7. Specifies the SNTP software version to use, and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3.

Menu: Viewing and Configuring SNTP

To View, Enable, and Modify SNTP Time Protocol:

1. From the Main Menu, select:

2. Switch Configuration...

1. System Information

```

===== CONSOLE - MANAGER MODE =====
Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes

Time Sync Method [None]: TIMEP ← Time Protocol Selection Parameter
TimeP Mode [Disabled] : Disabled      - TIMEP
                                      - SNTP
                                      - None

Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->  Cancel  Edit  Save  Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
  
```

Figure 9-1. The System Information Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **System Name** field.
3. Use **↓** to move the cursor to the **Time Sync Method** field.

4. Use the Space bar to select **SNTP**, then press **↓** once to display and move to the **SNTP Mode** field.
5. Do one of the following:
 - Use the Space bar to select the **Broadcast** mode, then press **↓** to move the cursor to the **Poll Interval** field, and go to step 6. (For Broadcast mode details, see “SNTP Operating Modes” on page 9-3.)

```
Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Broadcast
Poll Interval (sec) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

Figure 9-2. Time Configuration Fields for SNTP with Broadcast Mode

- Use the Space bar to select the **Unicast** mode, then do the following:
 - i. Press **→** to move the cursor to the **Server Address** field.
 - ii. Enter the IP address of the SNTP server you want the switch to use for time synchronization.

Note: This step replaces any previously configured server IP address. If you will be using backup SNTP servers (requires use of the CLI), then see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-25.

- iii. Press **↓** to move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step (step ii). If you are unsure which version to use, ProCurve recommends leaving this value at the default setting of **3** and testing SNTP operation to determine whether any change is necessary.

Note: Using the menu to enter the IP address for an SNTP server when the switch already has one or more SNTP servers configured causes the switch to delete the primary SNTP server from the server list and to select a new primary SNTP server from the IP address(es) in the updated list. For more on this topic, see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-25.

- iv. Press **→** to move the cursor to the **Poll Interval** field, then go to step 6.

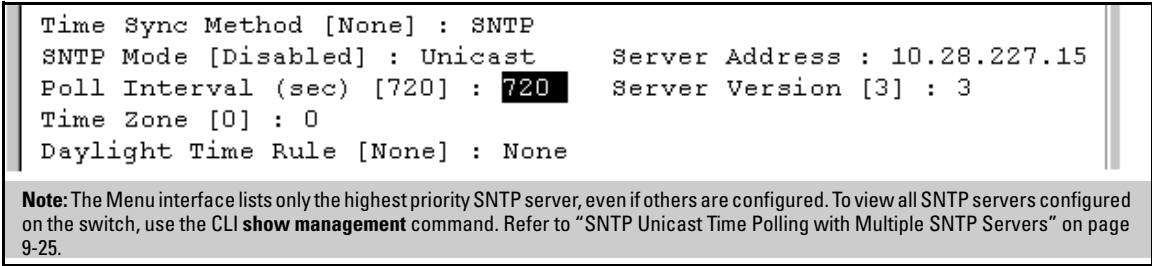


Figure 9-3. SNTP Configuration Fields for SNTP Configured with Unicast Mode

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see table 9-1, “SNTP Parameters”, on page 9-5.)
7. Press **[Enter]** to return to the Actions line, then **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring SNTP

CLI Commands Described in this Section

SNTP Command	Page
show sntp	9-9
[no] timesync	9-10 and ff., 9-14
sntp broadcast	9-11
sntp unicast	9-11
sntp server	9-11 and ff.
Protocol Version	9-13
poll-interval	9-13
no sntp	9-14

This section describes how to use the CLI to view, enable, and configure SNTP parameters.

Viewing the Current SNTP Configuration

Syntax: show sntp

*This command lists both the time synchronization method (**TimeP**, **SNTP**, or **None**) and the SNTP configuration, even if SNTP is not the selected time protocol.*

For example, if you configured the switch with SNTP as the time synchronization method, then enabled SNTP in broadcast mode with the default poll interval, show sntp lists the following:

```
ProCurve(config)# show sntp

SNTP Configuration

Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720

IP Address          Protocol Version
-----
10.10.28.101        3
10.11.35.117        3
10.12.115.86        3
```

Figure 9-4. Example of SNTP Configuration When SNTP Is the Selected Time Synchronization Method

In the factory-default configuration (where TimeP is the selected time synchronization method), **show sntp** still lists the SNTP configuration even though it is not currently in use. For example:

```
ProCurve(config)# show sntp

SNTP Configuration

Time Sync Mode: Timep
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720

IP Address          Protocol Version
-----
10.10.28.101        3
10.11.35.117        3
10.12.115.86        3
```

Even though, in this example, TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

Figure 9-5. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

Syntax: show management

This command can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.

```
ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : 10.10.28.100

SNTP Server Address Protocol Version
-----
10.10.28.101          3
10.11.35.117         3
10.12.115.86         3

Default Gateway      : 10.30.248.1

VLAN Name      MAC Address      | IP Address
-----
DEFAULT_VLAN  0004ea-5e2000    | 10.30.248.184
VLAN100       0004ea-5e2000    | 10.29.16.105
```

Figure 9-6. Example of Display Showing IP Addressing for All Configured Time Servers and VLANs

Configuring (Enabling or Disabling) the SNTP Mode

Enabling the SNTP mode means to configure it for either broadcast or unicast mode. Remember that to run SNTP as the switch's time synchronization protocol, you must also select SNTP as the time synchronization method by using the CLI **timesync** command (or the Menu interface **Time Sync Method** parameter).

Syntax: timesync sntp

Selects SNTP as the time protocol.

sntp < broadcast | unicast >

Enables the SNTP mode (below and page 9-11).

Syntax: sntp server < ip-addr >

Required only for unicast mode page 9-11).

Syntax: `sntp poll-interval < 30 - 720 >`

Enabling the SNTP mode also enables the SNTP poll interval (default: 720 seconds; page 9-13).

Enabling SNTP in Broadcast Mode. Because the switch provides an SNTP polling interval (default: 720 seconds), you need only these two commands for minimal SNTP broadcast configuration:

Syntax: `timesync sntp`

Selects SNTP as the time synchronization method.

Syntax: `sntp broadcast`

*Configures **broadcast** as the SNTP mode.*

For example, suppose:

- Time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method).
- You want to:
 1. View the current time synchronization.
 2. Select SNTP as the time synchronization mode.
 3. Enable SNTP for Broadcast mode.
 4. View the SNTP configuration again to verify the configuration.

The commands and output would appear as follows:

```
ProCurve(config)# show sntp ❶
SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
ProCurve(config)# timesync sntp ❷
ProCurve(config)# sntp broadcast ❸
ProCurve(config)# show sntp ❹
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

❶ `show sntp` displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.

❹ `show sntp` again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.

Figure 9-7. Example of Enabling SNTP Operation in Broadcast Mode

Enabling SNTP in Unicast Mode. Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for Unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up

to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing Unicast server with another. To add a second or third server, you must use the CLI. For more on SNTP operation with multiple servers, see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 25.

Syntax: `timesync sntp`

Selects SNTP as the time synchronization method.

Syntax: `sntp unicast`

Configures the SNTP mode for Unicast operation.

Syntax: `sntp server <ip-addr> [version]`

Specifies the SNTP server. The default server version is 3.

Syntax: `no sntp server <ip-addr>`

Deletes the specified SNTP server.

Note

Deleting an SNTP server when only one is configured disables SNTP unicast operation.

For example, to select SNTP and configure it with unicast mode and an SNTP server at 10.28.227.141 with the default server version (3) and default poll interval (720 seconds):

```
ProCurve(config)# timesync sntp
```

Selects SNTP.

```
ProCurve(config)# sntp unicast
```

Activates SNTP in Unicast mode.

```
ProCurve(config)# sntp server 10.28.227.141
```

Specifies the SNTP server and accepts the current SNTP server version (default: 3).

```
ProCurve(config)# show sntp
SNTP Configuration
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720
IP Address          Protocol Version
-----
10.28.227.141      3
```

In this example, the **Poll Interval** and the **Protocol Version** appear at their default settings.

Note: Protocol Version appears only when there is an IP address configured for an SNTP server.

Figure 9-8. Example of Configuring SNTP for Unicast Operation

If the SNTP server you specify uses SNTP version 4 or later, use the `sntp server` command to specify the correct version number. For example, suppose you learned that SNTP version 4 was in use on the server you specified above (IP address 10.28.227.141). You would use the following commands to delete the server IP address and then re-enter it with the correct version number for that server:

```
ProCurve(config)# no sntp server 10.28.227.141
ProCurve(config)# sntp server 10.28.227.141 4
ProCurve(config)# show sntp
SNTP Configuration
Time Sync Mode: Sntp
SNTP Mode : Broadcast
Poll Interval (sec) [720] : 600
IP Address          Protocol Version
-----
10.28.227.141      4
```

Deletes unicast SNTP server entry.

Re-enters the unicast server with a non-default protocol version.

show sntp displays the result.

Figure 9-9. Example of Specifying the SNTP Protocol Version Number

Changing the SNTP Poll Interval.

Syntax: `sntp poll-interval <30..720>`

Specifies how long the switch waits between time polling intervals. The default is 720 seconds and the range is 30 to 720 seconds. (This parameter is separate from the poll interval parameter used for Timep operation.)

For example, to change the poll interval to 300 seconds:

```
ProCurve(config)# sntp poll-interval 300
```

Disabling Time Synchronization Without Changing the SNTP

Configuration. The recommended method for disabling time synchronization is to use the **timesync** command.

Syntax: no timesync

Halts time synchronization without changing your SNTP configuration.

For example, suppose SNTP is running as the switch's time synchronization protocol, with **Broadcast** as the SNTP mode and the factory-default polling interval. You would halt time synchronization with this command:

```
ProCurve(config)# no timesync
```

If you then viewed the SNTP configuration, you would see the following:

```
ProCurve(config)# show sntp
SNTP Configuration
  Time Sync Mode: Disabled
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

Figure 9-10. Example of SNTP with Time Synchronization Disabled

Disabling the SNTP Mode. If you want to prevent SNTP from being used even if selected by **timesync** (or the Menu interface's **Time Sync Method** parameter), configure the SNTP mode as disabled.

Syntax: no sntp

*Disables SNTP by changing the SNTP mode configuration to **Disabled**.*

For example, if the switch is running SNTP in Unicast mode with an SNTP server at 10.28.227.141 and a server version of 3 (the default), **no sntp** changes the SNTP configuration as shown below, and disables time synchronization on the switch.

```
ProCurve(config)# no sntp
ProCurve(config)# show sntp
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
  IP Address          Protocol Version
  -----
  10.28.227.141      3
```

Even though the **Time Sync Mode** is set to **Sntp**, time synchronization is disabled because **no sntp** has disabled the **SNTP Mode** parameter.

Figure 9-11. Example of Disabling Time Synchronization by Disabling the SNTP Mode

TimeP: Viewing, Selecting, and Configuring

TimeP Feature	Default	Menu	CLI	Web
view the Timep time synchronization configuration	n/a	page 9-17	page 9-19	—
select Timep as the time synchronization method	TIMEP	page 9-15	pages 9-21 ff.	—
disable time synchronization	timep	page 9-17	page 9-23	—
enable the Timep mode	Disabled			—
DHCP	—	page 9-17	page 9-21	—
manual	—	page 9-18	page 9-22	—
none/disabled	—	page 9-17	page 9-23	—
change the SNTP poll interval	720 minutes	page 9-18	page 9-23	—

Table 9-2. Timep Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either TIMEP (the default), SNTP, or None as the time synchronization method.
Timep Mode	
Disabled	The Default. Timep does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI timesync command.
DHCP	When Timep is selected as the time synchronization method, the switch attempts to acquire a Timep server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the Timep poll interval. If the switch does not receive a Timep server IP address, it cannot perform time synchronization updates.
Manual	When Timep is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the Timep poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur.
Server Address	Used only when the TimeP Mode is set to Manual . Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server.
Poll Interval (minutes)	Default: 720 minutes. Specifies the interval the switch waits between attempts to poll the TimeP server for updates.

Menu: Viewing and Configuring TimeP

To View, Enable, and Modify the TimeP Protocol:

1. From the Main Menu, select:

2. Switch Configuration...

1. System Information

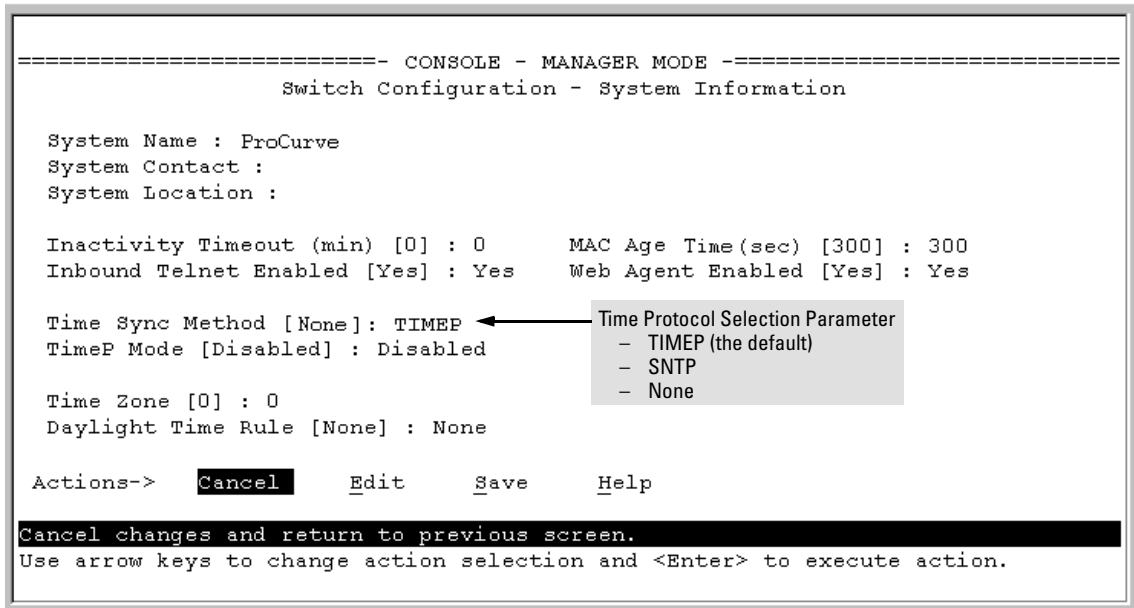







Figure 9-12. The System Information Screen (Default Values)

- Press [E] (for **E**dit). The cursor moves to the **System Name** field.
2. Use  to move the cursor to the **Time Sync Method** field.
 3. If **TIMEP** is not already selected, use the Space bar to select **TIMEP**, then press  once to display and move to the **TimeP Mode** field.
 4. Do one of the following:
 - Use the Space bar to select the **DHCP** mode, then press  to move the cursor to the **Poll Interval** field, and go to step 6.

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : DHCP
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

- Use the Space bar to select the **Manual** mode.
 - i. Press  to move the cursor to the **Server Address** field.
 - ii. Enter the IP address of the TimeP server you want the switch to use for time synchronization.
Note: This step replaces any previously configured TimeP server IP address.
 - iii. Press  to move the cursor to the **Poll Interval** field, then go to step 6.
- 5. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.

Press **[Enter]** to return to the Actions line, then **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring TimeP

CLI Commands Described in this Section

Command	Page
show timep	9-19
[no] timesync	9-20 ff., 9-23
ip timep	
dhcp	9-21
manual	9-22
server <ip-addr>	9-22
interval	9-23
no ip timep	9-23

This section describes how to use the CLI to view, enable, and configure TimeP parameters.

Viewing the Current TimeP Configuration

Using different **show** commands, you can display either the full TimeP configuration or a combined listing of all TimeP, SNTP, and VLAN IP addresses configured on the switch.

Syntax: show timep

*This command lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol. (If the TimeP Mode is set to **Disabled** or **DHCP**, then the **Server** field does not appear.)*

For example, if you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, **show timep** lists the following:

```
ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Timep
TimeP Mode [Disabled] : DHCP          Server Address : 10.10.28.100
Poll Interval (min) [720] : 720
```

Figure 9-13. Example of TimeP Configuration When TimeP Is the Selected Time Synchronization Method

If SNTP is the selected time synchronization method, **show timep** still lists the TimeP configuration even though it is not currently in use:

```
ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Sntp
TimeP Mode [Disabled] : Manual          Server Address : 10.10.28.100
Poll Interval (min) [720] : 720
```

Even though, in this example, SNTP is the current time synchronization method, the switch maintains the TimeP configuration.

Figure 9-14. Example of TimeP Configuration When TimeP Is Not the Selected Time Synchronization Method

Syntax: show management

This command can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.

```
ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : 10.10.28.100

SNTP Server Address Protocol Version
-----
10.10.28.101          3
10.11.35.117         3
10.12.115.86         3

Default Gateway      : 10.30.248.1

VLAN Name      MAC Address      | IP Address
-----
DEFAULT_VLAN  0004ea-5e2000      | 10.30.248.184
VLAN100       0004ea-5e2000      | 10.29.16.105
```

Figure 9-15. Example of Display Showing IP Addressing for All Configured Time Servers and VLANs

Configuring (Enabling or Disabling) the TimeP Mode

Enabling the TimeP mode means to configure it for either broadcast or unicast mode. Remember that to run TimeP as the switch's time synchronization protocol, you must also select TimeP as the time synchronization method by using the CLI `timesync` command (or the Menu interface **Time Sync Method** parameter).

Syntax: `timesync timep`
Selects TimeP as the time protocol.

Syntax: `ip timep < dhcp | manual >`
Enables the selected TimeP mode.

Syntax: `no ip timep`
Disables the TimeP mode.

Syntax: `no timesync`
Disables the time protocol.

Enabling TimeP in DHCP Mode. Because the switch provides a TimeP polling interval (default: 720 minutes), you need only these two commands for a minimal TimeP DHCP configuration:

Syntax: `timesync timep`

Selects TimeP as the time synchronization method.

Syntax: `ip timep dhcp`

Configures DHCP as the TimeP mode.

For example, suppose:

- Time synchronization is configured for SNTP.
- You want to:
 1. View the current time synchronization.
 2. Select TimeP as the time synchronization mode.
 3. Enable TimeP for DHCP mode.
 4. View the TimeP configuration.

The commands and output would appear as follows:

```
ProCurve(config)# show timep 1 show timep displays the TimeP configuration and also shows
Timep Configuration          that SNTP is the currently active time synchronization mode.
Time Sync Mode: Sntp
TimeP Mode : Disabled

ProCurve(config)# timesync timep 2

ProCurve(config)# ip timep dhcp 3

ProCurve(config)# show timep 4 show timep again displays the TimeP configuration and shows that TimeP is
Timep Configuration          now the currently active time synchronization mode.
Time Sync Mode: Timep
TimeP Mode : DHCP    Poll Interval (min) : 720
```

Figure 9-16. Example of Enabling TimeP Operation in DHCP Mode

Enabling Timep in Manual Mode. Like DHCP mode, configuring TimeP for **Manual** mode enables TimeP. However, for manual operation, you must also specify the IP address of the TimeP server. (The switch allows only one TimeP server.) To enable the TimeP protocol:

Syntax: timesync timep

Selects Timep.

Syntax: ip timep manual < ip-addr >

Activates TimeP in Manual mode with a specified TimeP server.

Syntax: no ip timep

Disables TimeP.

Note

To change from one TimeP server to another, you must (1) use the **no ip timep** command to disable TimeP mode, and then reconfigure TimeP in Manual mode with the new server IP address.

For example, to select TimeP and configure it for manual operation using a TimeP server address of 10.28.227.141 and the default poll interval (720 minutes, assuming the TimeP poll interval is already set to the default):

```
ProCurve(config)# timesync timep
```

Selects TimeP.

```
ProCurve(config)# ip timep manual 10.28.227.141
```

Activates TimeP in Manual mode.

```
ProCurve(config)# timesync timep
ProCurve(config)# ip timep manual 10.28.227.141

ProCurve(config)# show timep
Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Manual          Server Address : 10.28.227.141
  Poll Interval (min) : 720
```

Figure 9-17. Example of Configuring Timep for Manual Operation

Changing the TimeP Poll Interval. This command lets you specify how long the switch waits between time polling intervals. The default is 720 minutes and the range is 1 to 9999 minutes. (This parameter is separate from the poll interval parameter used for SNTP operation.)

Syntax: `ip timep < dhcp | manual > interval < 1 - 9999 >`

For example, to change the poll interval to 60 minutes:

```
ProCurve(config)# ip timep interval 60
```

Disabling Time Synchronization Without Changing the TimeP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your TimeP configuration.

Syntax: `no timesync`

*Disables time synchronization by changing the **Time Sync Mode** configuration to **Disabled**.*

For example, suppose TimeP is running as the switch's time synchronization protocol, with **DHCP** as the TimeP mode, and the factory-default polling interval. You would halt time synchronization with this command:

```
HPswitch(config)# no timesync
```

If you then viewed the TimeP configuration, you would see the following:

```
ProCurve(config)# show timep
Timep Configuration
  Time Sync Mode: Disabled
  TimeP Mode : DHCP      Poll Interval (min) : 720
```

Figure 9-18. Example of TimeP with Time Synchronization Disabled

Disabling the TimeP Mode. Disabling the TimeP mode means to configure it as disabled. (Disabling TimeP prevents the switch from using it as the time synchronization protocol, even if it is the selected **Time Sync Method** option.)

Syntax: `no ip timep`

*Disables TimeP by changing the TimeP mode configuration to **Disabled**.*

Time Protocols

TimeP: Viewing, Selecting, and Configuring

For example, if the switch is running TimeP in DHCP mode, **no ip timep** changes the TimeP configuration as shown below, and disables time synchronization.

```
ProCurve(config)# no ip timep

ProCurve(config)# show timep
Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Disabled
```

Even though the Time Sync Mode is set to Timep, time synchronization is disabled because no ip timep has disabled the TimeP Mode parameter.

Figure 9-19. Example of Disabling Time Synchronization by Disabling the TimeP Mode Parameter

SNTP Unicast Time Polling with Multiple SNTP Servers

When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured with either the Server Address parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured **Poll Interval** time has expired.

Address Prioritization

If you use the CLI to configure multiple SNTP servers, the switch prioritizes them according to the decimal values of their IP addresses. That is, the switch compares the decimal value of the octets in the addresses and orders them accordingly, with the lowest decimal value assigned as the primary address, the second-lowest decimal value assigned as the next address, and the third-lowest decimal value as the last address. If the first octet is the same between two of the addresses, the second octet is compared, and so on. For example:

SNTP Server IP Address	Server Ranking According to Decimal Value of IP Address
10.28.227.141	Primary
10.28.227.153	Secondary
10.29.227.100	Tertiary

Displaying All SNTP Server Addresses Configured on the Switch

The System Information screen in the menu interface displays only one SNTP server address, even if the switch is configured for two or three servers. The CLI **show management** command displays all configured SNTP servers on the switch.

```
ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : Disabled

SNTP Server Address Protocol Version
-----
10.28.227.141      3
10.28.227.153      3
10.29.227.100      3

Default Gateway    : 15.30.248.1

VLAN Name      MAC Address      | IP Address
-----
DEFAULT_VLAN   0004ea-5e2000    | 15.30.248.184
VLAN28         0004ea-5e2000    | 10.28.227.100
VLAN29         0004ea-5e2000    | 10.29.227.53
```

Figure 9-20. Example of How To List All SNTP Servers Configured on the Switch

Adding and Deleting SNTP Server Addresses

Adding Addresses. As mentioned earlier, you can configure one SNTP server address using either the Menu interface or the CLI. To configure a second and third address, you must use the CLI. For example, suppose you have already configured the primary address in the above table (10.28.227.141). To configure the remaining two addresses, you would do the following:


```
ProCurve(config)# sntp server 10.29.227.100
ProCurve(config)# sntp server 10.28.227.153
ProCurve(config)# show sntp
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
  IP Address      Protocol Version
  -----
  10.28.227.141   3
  10.28.227.153   3
  10.29.227.100   3
```

Prioritized list of SNTP Server IP Addresses




Figure 9-21. Example of SNTP Server Address Prioritization

Note

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Deleting Addresses. To delete an address, you must use the CLI. If there are multiple addresses and you delete one of them, the switch re-orders the address priority. (See “Address Prioritization” on page 25.)

Syntax: no sntp server < ip-addr >

For example, to delete the primary address in the above example (and automatically convert the secondary address to primary):

```
ProCurve(config)# no sntp server 10.28.227.141
```

Menu: Operation with Multiple SNTP Server Addresses Configured

When you use the Menu interface to configure an SNTP server IP address, the new address writes over the current primary address, if one is configured. If there are multiple addresses configured, the switch re-orders the addresses according to the criteria described under “Address Prioritization” on page 25. For example, suppose the switch already has the following three SNTP server IP addresses configured.

- 10.28.227.141 (primary)
- 10.28.227.153 (secondary)
- 10.29.227.100 (tertiary)

If you use the Menu interface to add 10.28.227.160, the new prioritized list will be:

New Address List	Address Status
10.28.227.153	New Primary (The former primary, 10.28.227.141 was deleted when you used the menu to add 10.28.227.160.)
10.28.227.160	New Secondary
10.29.227.100	Same Tertiary (This address still has the highest decimal value.)

SNTP Messages in the Event Log

If an SNTP time change of more than three seconds occurs, the switch’s event log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Port Status and Basic Configuration

Contents

Overview	10-2
Viewing Port Status and Configuring Port Parameters	10-2
Menu: Port Configuration	10-6
CLI: Viewing Port Status and Configuring Port Parameters	10-8
Using the CLI To Enable or Disable Ports and Configure Port Mode	10-9
Enabling or Disabling Flow Control	10-11
Configuring a Broadcast Limit on the Switch	10-14
Configuring Auto-MDIX	10-15
Web: Viewing Port Status and Configuring Port Parameters	10-18
Using Friendly (Optional) Port Names	10-18
Configuring and Operating Rules for Friendly Port Names	10-19
Configuring Friendly Port Names	10-19
Displaying Friendly Port Names with Other Port Data	10-21

Overview

This chapter describes how to view the current port configuration and how to configure ports to non-default settings, including

- Enable/Disable
 - Mode (speed and duplex)
 - Flow Control
 - Broadcast Limit
-

Viewing Port Status and Configuring Port Parameters

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port status	n/a	page 10-6	page 10-8	page 10-18
configuring ports	Refer to Table 10-1 on pages 10-3 thru 10-5	page 10-7	page 10-9	page 10-18
configuring auto-mdix			page 9-11	

Note On Connecting Transceivers to Fixed-Configuration Devices

If the switch either fails to show a link between an installed transceiver and another device, or demonstrates errors or other unexpected behavior on the link, check the port configuration on both devices for a speed and/or duplex (mode) mismatch. To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface (page 10-6) or **show interfaces brief** in the CLI (page 10-8).

Table 10-1. Status and Parameters for Each Port Type

Status or Parameter	Description
Enabled	<p>Yes (default): The port is ready for a network connection.</p> <p>No: The port will not operate, even if properly connected in a network. Use this setting, for example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes.</p>
Status (read-only)	<p>Up: The port senses a link beat.</p> <p>Down: The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, refer to the installation manual you received with the switch. Refer also to appendix C, “Troubleshooting” (in this manual).</p>
Mode	<p>The port’s speed and duplex (data transfer operation) setting.</p> <p>10/100Base-T Ports:</p> <ul style="list-style-type: none"> • auto-mdix (default): Senses speed and negotiates with the port at the other end of the link for data transfer operation (half-duplex or full-duplex). <ul style="list-style-type: none"> Note: Ensure that the device attached to the port is configured for the same setting that you select here. Also, if “Auto” is used, the device to which the port is connected must operate in compliance with the IEEE 802.3u “Auto Negotiation” standard for 100Base-T networks. If the other device does not comply with the 802.3u standard, or is not set to Auto, then the port configuration on the switch must be manually set to match the port configuration on the other device. To see what the switch negotiates for the Auto setting, use the CLI show interfaces brief command or the “3. Port Status” option under “1. Status and Counters” in the menu interface. • mdi: Sets the port to connect with a PC using a crossover cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables) • mdix: Sets the port to connect with a PC using a straight-through cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables) • Auto-10: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). ProCurve recommends Auto-10 for links between 10/100 auto-sensing ports connected with Cat3 cabling. (Cat5 cabling is required for 100 Mbps links.). • 10HDx: 10 Mbps, Half-Duplex • 10FDx: 10 Mbps, Full-Duplex • 100HDx: 100 Mbps, Half-Duplex • 100FDx: 100 Mbps, Full-Duplex <p>100FX Ports:</p> <ul style="list-style-type: none"> • 100HDx: 100 Mbps, Half-Duplex • 100FDx (default): 100 Mbps, Full-Duplex

— Continued —

Status or Parameter	Description
---------------------	-------------

— Continued From Previous Page —

100/1000Base-T Ports:

- **auto-mdix** (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI).
To see what the switch negotiates for the Auto setting, use the CLI **show interfaces brief** command or the “**3. Port Status**” option under “**1. Status and Counters**” in the menu interface.
- **mdi**: Sets the port to connect with a PC using a crossover cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **mdix**: Sets the port to connect with a PC using a straight-through cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **Auto-100**: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **Auto-1000**: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **100Hdx**: Uses 100 Mbps, half-duplex.
- **100Fdx**: Uses 100 Mbps, Full-Duplex

Notes:

- Changing the port speed on a transceiver port requires a reboot of the switch.
- Ensure that the device attached to the port is configured for the same setting that you select here. Also, if “Auto” is used, the device to which the port connects must also be configured to “Auto” and operate in compliance with the IEEE 802.3ab “Auto Negotiation” standard for 1000Base-T networks.

10/100/1000Base-T Ports:

- **auto-mdix** (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI).
To see what the switch negotiates for the Auto setting, use the CLI **show interfaces brief** command or the “**3. Port Status**” option under “**1. Status and Counters**” in the menu interface.
- **mdi**: Sets the port to connect with a PC using a crossover cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **mdix**: Sets the port to connect with a PC using a straight-through cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **Auto-10**: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). ProCurve recommends Auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.).
- **10HDx**: 10 Mbps, Half-Duplex
- **10FDx**: 10 Mbps, Full-Duplex
- **Auto-100**: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **Auto-1000**: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **100Hdx**: Uses 100 Mbps, half-duplex.
- **100Fdx**: Uses 100 Mbps, Full-Duplex

— Continued on Next Page —

Status or Parameter	Description
<i>— Continued From Previous Page —</i>	
	<p>Gigabit Fiber-Optic Ports (Gigabit-SX, Gigabit-LX, and Gigabit-LH):</p> <ul style="list-style-type: none"> • 1000FDx: 1000 Mbps (1 Gbps), Full Duplex only • Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port.
	<p>10-Gigabit CX4 Copper Ports:</p> <ul style="list-style-type: none"> • Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed.
	<p>10-Gigabit SC Fiber-Optic Ports:</p> <ul style="list-style-type: none"> • Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed.
Auto-MDIX	<p>The switch supports Auto-MDIX on 10Mb, 100Mb, and 1 Gb T/TX (copper) ports. (Fiber ports and 10-gigabit ports do not use this feature.)</p> <ul style="list-style-type: none"> • Automdix: Configures the port for automatic detection of the cable type (straight-through or crossover). • MDI: Configures the port to connect to a switch, hub, or other MDI-X device with a straight-through cable. • MDIX: Configures the port to connect to a PC or other MDI device with a straight-through cable.
Flow Control	<ul style="list-style-type: none"> • Disabled (default): The port does not generate flow control packets, and drops any flow control packets it receives. • Enabled: The port uses 802.3x Link Layer Flow Control, generates flow control packets, and processes received flow control packets. <p>With the port mode set to Auto (the default) and Flow Control enabled, the switch negotiates Flow Control on the indicated port. If the port mode is not set to Auto, or if Flow Control is disabled on the port, then Flow Control is not used. You must enable flow control globally on the switch before enabling it on individual ports. Also, you must disable flow control on the individual ports before disabling it globally on the switch. Note that flow control must be enabled on both ends of a link.</p>
Broadcast Limit	<p>Specifies the percentage of the theoretical maximum network bandwidth that can be used for broadcast and multicast traffic. Any broadcast or multicast traffic exceeding that limit will be dropped. Zero (0) means the feature is disabled.</p> <p>Series 5300xl and Series 4200vl Switches: The broadcast-limit command operates at the global configuration context level to set the broadcast limit for all ports on the switch.</p> <p>Series 3400cl and Series 6400cl Switches: The broadcast-limit command operates at the port context level to set the broadcast limit on a per-port basis.</p> <p>Note: This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.</p>

Menu: Port Configuration

From the menu interface, you can view and change the port configuration.

Using the Menu To View Port Configuration. The menu interface displays the configuration for ports and (if configured) any trunk groups.

From the Main Menu, select:

1. Status and Counters ...

3. Port Status (3400cl and 6400cl switches)

— or —

4. Port Status (5300cl and 4200vl switches)

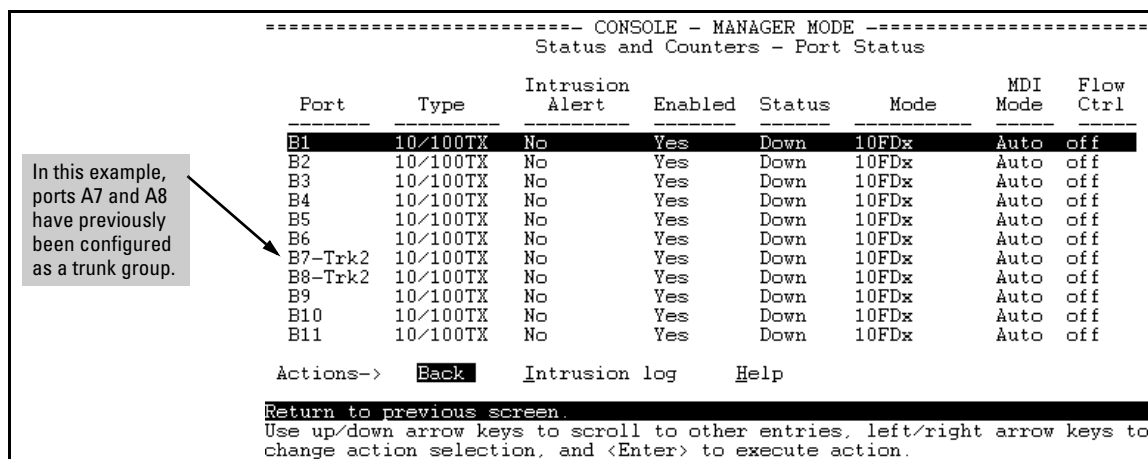


Figure 10-1. Example of a 5300xl Port Status Screen

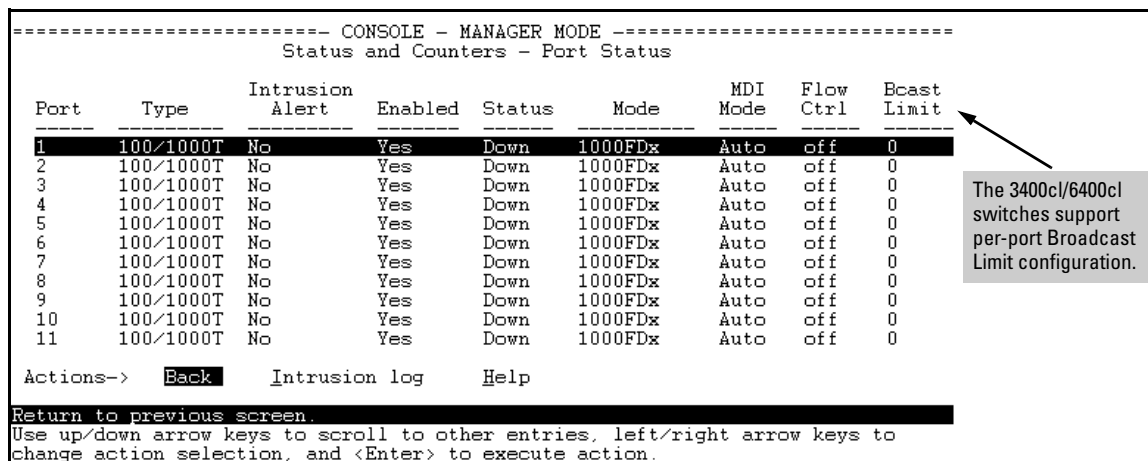


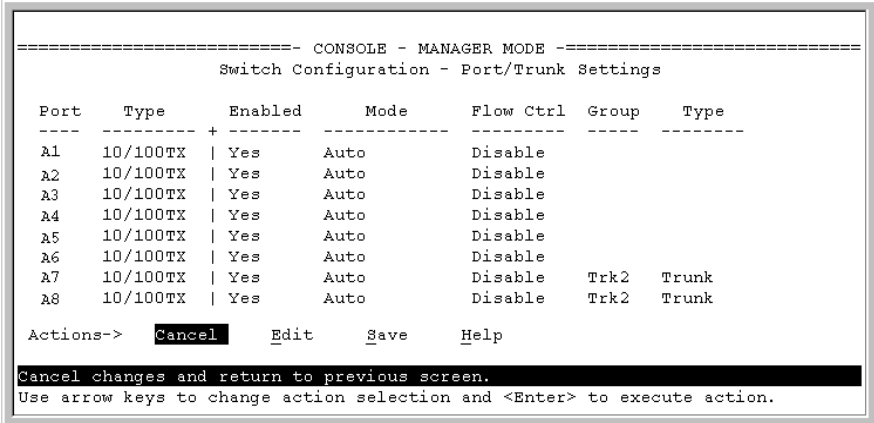
Figure 10-2. Example of a 3400cl Port Status Screen

Using the Menu To Configure Ports.

Note

The menu interface uses the same screen for configuring both individual ports and port trunk groups. For information on port trunk groups, refer to chapter 13, “Port Trunking” .

1. From the Main Menu, Select:
2. Switch Configuration...
2. Port/Trunk Settings



```
=====  CONSOLE - MANAGER MODE  =====
                        Switch Configuration - Port/Trunk Settings

Port   Type      Enabled  Mode      Flow Ctrl  Group   Type
-----+-----+-----+-----+-----+-----+-----
A1     10/100TX  | Yes    Auto      Disable
A2     10/100TX  | Yes    Auto      Disable
A3     10/100TX  | Yes    Auto      Disable
A4     10/100TX  | Yes    Auto      Disable
A5     10/100TX  | Yes    Auto      Disable
A6     10/100TX  | Yes    Auto      Disable
A7     10/100TX  | Yes    Auto      Disable    Trk2    Trunk
A8     10/100TX  | Yes    Auto      Disable    Trk2    Trunk

Actions->  Cancel  Edit    Save    Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 10-3. Example of Port/Trunk Settings with a Trunk Group Configured

2. Press [E] (for Edit). The cursor moves to the **Enabled** field for the first port.
3. Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for Save).

CLI: Viewing Port Status and Configuring Port Parameters

Port Status and Configuration Commands

show interfaces brief	page 10-9
show interfaces config	page 10-9
interface	page 10-9
disable/enable	page 10-9
speed-duplex	page 10-9
flow-control	page 10-11
broadcast-limit	page 10-14
auto-mdix	page 10-15

From the CLI, you can configure and view all port parameter settings and view all port status indicators.

Using the CLI To View Port Status. Use the following commands to display port status and configuration.

Syntax: show interfaces [brief | config | < port-list >]

brief: Lists the current operating status for all ports on the switch.

config: Lists a subset of configuration data for all ports on the switch; that is, for each port, the display shows whether the port is enabled, the operating mode, and whether it is configured for flow control.

< port-list >: Shows a summary of network traffic handled by the specified ports.

*Series 3400cl and Series 6400cl switches include per-port broadcast limit settings in the **show interfaces** and **show interfaces brief** display outputs.*

The next two figures list examples of the output of the above two command options for the same port configuration.

ProCurve(config)# show interfaces brief						This screen shows current port operating status.		
Status and Counters - Port Status						Note: The (per-port) Bcast Limit column appears only on the 3400cl and 6400cl switches. (The 5300xl switches apply a global broadcast limit)		
Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl	Bcast Limit
1	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
2	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
3	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
4	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
5	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
6	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
24	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
25	10GbE-CX4	No	Yes	Down	10-Gig	n/a	off	0
26	10GbE-LR	No	Yes	Down	10-Gig	n/a	off	0

3400cl/
6400cl
Switches
Only

Figure 10-4. Example of a Show Interfaces Brief Command Listing

ProCurve(config)# show interface config						This screen shows current port configuration.	
Port Settings							
Port	Type	Enabled	Mode	Flow Ctrl	MDI		
1	100/1000T	Yes	Auto	Disable	Auto		
2	100/1000T	Yes	Auto	Disable	Auto		
3	100/1000T	Yes	Auto	Disable	Auto		
4	100/1000T	Yes	Auto	Disable	Auto		
5	100/1000T	Yes	Auto	Disable	Auto		
6	100/1000T	Yes	Auto	Disable	Auto		
⋮	⋮	⋮	⋮	⋮	⋮		
24	100/1000T	Yes	Auto	Disable	Auto		
25	10GbE-CX4	Yes	Auto	Disable			
26	10GbE-LR	Yes	Auto	Disable			

Figure 10-5. Example of a Show Interfaces Config Command Listing

Using the CLI To Enable or Disable Ports and Configure Port Mode

You can configure one or more of the following port parameters. For details, refer to table 10-1 on pages 10-3 thru 10-5.

Syntax: [no] interface <port-list>

[<disable | enable>]

*Disables or enables the port for network traffic. Does not use the **no** form of the command. (Default: **enable**.)*

[speed-duplex <auto-10|10-full|10-half|100-full|100-half|auto|auto-100|1000-full>]

*Specifies the port's data transfer speed and mode. Does not use the **no** form of the command. (Default: **auto**.)*

Note that in the above syntax you can substitute an “**int**” for “**interface**”; that is: **int < port-list >**.

For example, to configure ports C1 through C3 and port C6 for 100Mbps full-duplex, you would enter these commands:

```
ProCurve(config)# int c1-c3,c6 speed-duplex 100-full
```

Similarly, to configure a single port with the above command settings, you could either enter the same command with only the one port identified, or go to the *context level* for that port and then enter the command. For example, to enter the context level for port C6 and then configure that port for 100FDx:

```
ProCurve(config)# int e c6
ProCurve(eth-C6)# speed-duplex 100-full
```

If port C8 was disabled, and you wanted to enable it and configure it for 100FDx with flow-control active, you could do so with either of the following command sets. (Note that to enable flow control on individual ports, you must first enable it globally, as shown in these examples.)

The diagram illustrates two methods for configuring port C8. The first method (top) shows a sequence of commands in the configuration mode: enabling the port, setting speed-duplex to 100-full, enabling flow-control, and then selecting the port context. The second method (bottom) shows selecting the port context first, then enabling the port, setting speed-duplex to 100-full, and enabling flow-control. Callouts explain that the first set of commands enables and configures port C8 from the config level, while the second set selects the port context level and then applies the subsequent configuration commands to port C8. A third callout explains that the flow-control command in the first method enables flow control globally on the switch, which is required before enabling flow control on specific ports.

```
ProCurve(config)# int c8 enable
ProCurve(config)# int c8 speed-duplex 100-full
ProCurve(config)# flow-control
ProCurve(config)# int c8 flow-control

ProCurve(config)# flow-control

ProCurve(config)# int c8
ProCurve(eth-C8)# enable
ProCurve(eth-C8)# speed-duplex 100-full
ProCurve(eth-C8)# flow-control
```

These commands enable and configure port C8 from the config level:

This command enables flow control globally on the switch (which is required before you can enable flow control on specific ports).

These commands select the port C8 context level and then apply the subsequent configuration commands to port C8:

Figure 10-6. Examples of Two Methods for Changing a Port Configuration

Refer to “Enabling or Disabling Flow Control” on page 10-11 for more on flow control.

Enabling or Disabling Flow Control

- **3400cl/6400cl Switches:** Flow-Control on these switches is enabled and disabled on a per-port basis.
- **5300xl and 4200vl Switches:** You must first enable flow-control globally on the switch, and then enable it on the desired ports.

Note

You must enable flow control on both ports in a given link. Otherwise, flow control does not operate on the link, and appears as **Off** in the **show interfaces brief** port listing, even if flow control is configured as enabled on the port in the switch. (Refer to figure 10-4 on page 10-9.) Also, the port (speed-duplex) mode must be set to **Auto** (the default).

Flow Control on the 3400cl/6400cl Switches.

Syntax: [no] interface < port-list > flow-control

Enables or disables flow control packets on < port-list > ports. The “no” form of the command disables flow control. (Default: Disabled.)

Use the show interfaces brief command (figure 10-4 on page 10-9) to view the current per-port flow-control configuration.

Flow Control on the 5300xl and 4200vl Switches. As mentioned earlier, flow control operates on individual 5300xl and 4200vl switch ports after you first enable global flow control and then per-port flow control. The reverse is true for disabling flow control on all ports. (Disable per-port flow control, and then disable global flow control.) To disable flow control on some ports, while leaving it enabled on other ports, just disable it on the individual ports you want to exclude.

Syntax: [no] flow-control

Enables or disables flow-control globally on the switch, and is required before you can enable flow control on specific ports. To use the no form of the command to disable global flow-control, you must first disable flow-control on all ports on the switch. (Default: Disabled)

[no] interface < port-list > flow-control

Enables or disables flow control packets on the port. The “no” form of the command disables flow control on the individual ports. (Default: Disabled.)

Port Status and Basic Configuration

Viewing Port Status and Configuring Port Parameters

For example, suppose that:

1. You want to enable flow control on ports A1-A6.
2. Later, you decide to disable flow control on ports A5 and A6.
3. As a final step, you want to disable flow control on all ports.

Assuming that flow control is currently disabled on the switch, you would use these commands:

```
ProCurve(config)# flow-control
ProCurve(config)# int a1-a6 flow-control
ProCurve(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl
A1	10/100TX	No	Yes	Up	10FDx	on
A2	10/100TX	No	Yes	Up	10FDx	on
A3	10/100TX	No	Yes	Up	10FDx	on
A4	10/100TX	No	Yes	Up	10FDx	on
A5	10/100TX	No	Yes	Up	10FDx	on
A6	10/100TX	No	Yes	Up	10FDx	on
A7	10/100TX	No	Yes	Down	10HDx	off
A8	10/100TX	No	Yes	Up	10FDx	off
.						
.						
.						

Figure 10-7. Example of Configuring Flow Control for a Series of Ports

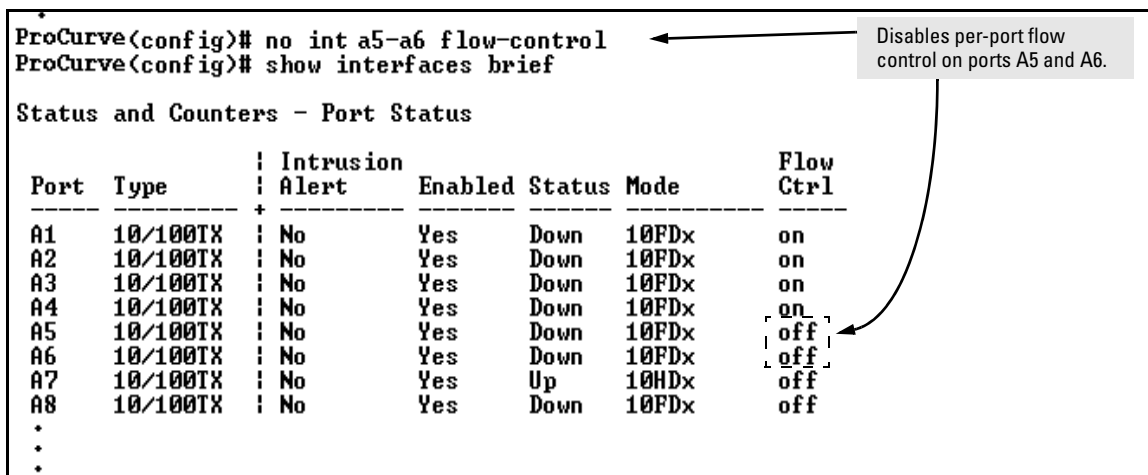


Figure 10-8. Example Continued from Figure 10-7

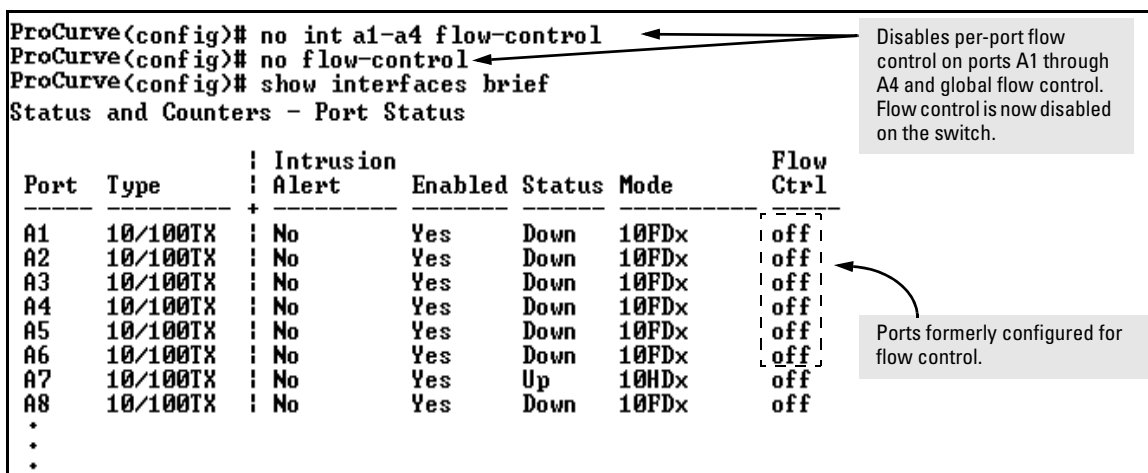


Figure 10-9. Example Continued from Figure 10-8

Configuring a Broadcast Limit on the Switch

- **3400cl/6400cl Switches:** Broadcast-Limit on these switches is configured as a percentage on a per-port basis.
- **5300xl and 4200vl Switches:** Broadcast-Limit on these switches is configured globally (on all ports) as a fixed limit.

Broadcast-Limit on the 3400cl/6400cl Switches.

Syntax: interface < port-list > broadcast-limit < 0 - 99 > (3400cl/6400cl Switches)

Configures the theoretical maximum bandwidth percentage that the specified switch ports use to limit broadcasts and multicasts. The switch drops any broadcast or multicast traffic exceeding that limit. Zero (0) disables the feature.

Note: This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

For example, to configure a broadcast limit of 70% on ports 1-12:

```
ProCurve(config)# interface 1-12 broadcast-limit 70
```

To later disable broadcast limiting on ports 11 and 12:

```
ProCurve(config)# interface 11-12 broadcast-limit 0
```

Broadcast-Limit on the 5300xl and 4200vl Switches.

Syntax: [no] broadcast-limit

Enables or disables broadcast limiting for outbound broadcasts on all ports on the switch. When enabled, this command limits outbound broadcast packets to 1,000 per second on each port, regardless of packet size.

Note: This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

Syntax: show config

Displays the startup-config file. The broadcast limit setting appears here if enabled and saved to the startup-config file.

Syntax: show running-config

Displays the running-config file. The broadcast limit setting appears here if enabled. If the setting is not also saved to the startup-config file, rebooting the switch returns broadcast limit to the setting currently in the startup-config file.

For example, the following command enables broadcast limiting on all ports on the switch:

```
ProCurve(config)# broadcast-limit
```

Configuring Auto-MDIX

Copper ports on the switch can automatically detect the type of cable configuration (MDI or MDI-X) on a connected device and adjust to operate appropriately.

This means you can use a “straight-through” twisted-pair cable or a “cross-over” twisted-pair cable for any of the connections—the port makes the necessary adjustments to accommodate either one for correct operation. The following port types on your switch support the IEEE 802.3ab standard, which includes the “Auto MDI/MDI-X” feature:

Series 5300xl and Series 4200vl Switches	Series 3400cl Switches
10/100-TX xl module ports	10/100/1000-T ports
100/1000-T xl module ports	
10/100/1000-T xl module ports	

(MDI/MDI-X does not apply to the optional 10-gigabit ports on the Series 3400cl switches or the 10-gigabit ports on the Series 6400cl switches.)

Using the above ports:

- If you connect a copper port using a straight-through cable on a Series 5300xl, Series 4200vl, or Series 3400cl switch to a port on another switch or hub that uses MDI-X ports, the 5300xl, 4200vl, or 3400cl switch port automatically operates as an MDI port.
- If you connect a copper port using a straight-through cable on a Series 5300xl, Series 4200vl, or Series 3400cl switch to a port on an end node, such as a server or PC, that uses MDI ports, the 5300xl, 4200vl, or 3400cl switch port automatically operates as an MDI-X port.

HP Auto-MDIX was developed for auto-negotiating devices, and was shared with the IEEE for the development of the IEEE 802.3ab standard. HP Auto-MDIX and the IEEE 802.3ab Auto MDI/MID-X feature are completely compatible. Additionally, HP Auto-MDIX supports operation in forced speed and duplex modes.

If you want more information on this subject please refer to the *IEEE 802.3ab Standard Reference*.

For more information on MDI-X, refer to the appendix titled “Switch Ports and Network Cables” in the *Installation and Getting Started Guide* for your switch.

Manual Override. If you require control over the MDI/MDI-X feature you can set the switch to either of two non-default modes:

- Manual MDI
- Manual MDI-X

Table 10-2 shows the cabling requirements for the MDI/MDI-X settings.

Table 10-2. Cable Types for Auto and Manual MDI/MDI-X Settings

Setting	MDI/MDI-X Device Type	
	PC or Other MDI Device Type	Switch, Hub, or Other MDI-X Device
Manual MDI	Crossover Cable	Straight-Through Cable
Manual MDI-X	Straight-Through Cable	Crossover Cable
Auto-MDI-X (The Default)	Either Crossover or Straight-Through Cable	

The Auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables.

Syntax: interface < port-list > mdix-mode < auto-mdix | mdi | mdix >

auto-mdix is the automatic, default setting. This configures the port for automatic detection of the cable (either straight-through or crossover).

mdi is the manual mode setting that configures the port for connecting to either a PC or other MDI device with a crossover cable, or to a switch, hub, or other MDI-X device with a straight-through cable.

mdix is the manual mode setting that configures the port for connecting to either a switch, hub, or other MDI-X device with a crossover cable, or to a PC or other MDI device with a straight-through cable.

Syntax: show interfaces config

Lists the current per-port Auto/MDI/MDI-X configuration.

Syntax: show interfaces brief

*Where a port is linked to another device, this command lists the MDI mode the port is currently using. In the case of ports configured for **Auto (auto-mdix)**, the MDI mode appears as either **MDI** or **MDIX**, depending upon which option the port has negotiated with the device on the other end of the link. In the case of ports configured for **MDI** or **MDIX**, the mode listed in this display matches the configured setting. If the link to another device was up, but has gone down, this command shows the last operating MDI mode the port was using. If a port on a given switch has not detected a link to another device since the last reboot, this command lists the MDI mode to which the port is currently configured.*

For example, **show interfaces config** displays the following data when port A1 is configured for **auto-mdix**, port A2 is configured for **mdi**, and port A3 is configured for **mdix**.

ProCurve(config)# show interfaces config						Per-Port MDI Configuration
Port Settings						
Port	Type	Enabled	Mode	Flow Ctrl	MDI	
A1	10/100TX	Yes	Auto	Disable	Auto	
A2	10/100TX	Yes	Auto	Disable	MDI	
A3	10/100TX	Yes	Auto	Disable	MDIX	
A4	10/100TX	Yes	Auto	Disable	Auto	
A5	10/100TX	Yes	Auto	Disable	Auto	
:	:	:	:	:	:	
:	:	:	:	:	:	

Figure 10-10. Example of Displaying the Current MDI Configuration

ProCurve(config)# show interfaces brief								Per-Port MDI Operating Mode
Status and Counters - Port Status								
Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl	
A1	10/100TX	No	Yes	Up	100FDx	MDIX	off	
A2	10/100TX	No	Yes	Up	100FDx	MDI	off	
A3	10/100TX	No	Yes	Up	100FDx	MDIX	off	
A4	10/100TX	No	Yes	Down	10FDx	Auto	off	
A5	10/100TX	No	Yes	Down	10FDx	Auto	off	
:	:	:	:	:	:	:	:	
:	:	:	:	:	:	:	:	

Figure 10-11. Example of Displaying the Current MDI Operating Mode

Note

Upgrading the Switch Series 5300xl Operating System from E_07.XX or earlier:

1. Copper ports in auto-negotiation still default to **auto-mdix** mode.
2. Copper ports in forced speed/duplex default to **mdix** mode.

For a fresh installation of the operating system, **auto-mdix** is the default.

Web: Viewing Port Status and Configuring Port Parameters

In the web browser interface:

1. Click on the **Configuration** tab.
2. Click on **[Port Configuration]**.
3. Select the ports you want to modify and click on **[Modify Selected Ports]**.
4. After you make the desired changes, click on **[Apply Settings]**.

Note that the web browser interface displays an existing port trunk group. However, to configure a port trunk group, you must use the CLI or the menu interface. For more on this topic, refer to chapter 13, “Port Trunking” .

Using Friendly (Optional) Port Names

Feature	Default	Menu	CLI	Web
Configure Friendly Port Names	Standard Port Numbering	n/a	page 19	n/a
Display Friendly Port Names	n/a	n/a	page 21	n/a

This feature enables you to assign alphanumeric port names of your choosing to augment automatically assigned numeric port names. This means you can configure meaningful port names to make it easier to identify the source of information listed by some **Show** commands. (Note that this feature *augments* port numbering, but *does not replace* it.)

Configuring and Operating Rules for Friendly Port Names

- At either the global or context configuration level you can assign a unique name to a port. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the **show name** [*port-list*], **show config**, and **show interface** <*port-number*> commands. They do not appear in the output of other show commands or in Menu interface screens. (See “Displaying Friendly Port Names with Other Port Data” on page 10-21.)
- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.
- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.
- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the **write memory** command.)

Configuring Friendly Port Names

Syntax: interface < *port-list* > name < *port-name-string* >
Assigns a port name to port-list.

Syntax: no interface < *port-list* > name
Deletes the port name from port-list.

Configuring a Single Port Name. Suppose that you have connected port A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

```
ProCurve(config)# int A3 name Bill_Smith@10.25.101.73
ProCurve(config)# write mem
ProCurve(config)# show name A3

Port Names
Port : A3
Type : 10/100TX
Name : Bill_Smith@10.25.101.73
```

Figure 10-12. Example of Configuring a Friendly Port Name

Configuring the Same Name for Multiple Ports. Suppose that you want to use ports A5 through A8 as a trunked link to a server used by a drafting group. In this case you might configure ports A5 through A8 with the name "Draft-Server:Trunk".

```
ProCurve(config)# int A5-A8 name Draft-Server:Trunk
ProCurve(config)# write mem
ProCurve(config)# show name 5-8

Port Names

Port : A5
Type : 10/100TX
Name : Draft-Server:Trunk

Port : A6
Type : 10/100TX
Name : Draft-Server:Trunk

Port : A7
Type : 10/100TX
Name : Draft-Server:Trunk

Port : A8
Type : 10/100TX
Name : Draft-Server:Trunk
```

Figure 10-13. Example of Configuring One Friendly Port Name on Multiple Ports

Displaying Friendly Port Names with Other Port Data

You can display friendly port name data in the following combinations:

- **show name:** Displays a listing of port numbers with their corresponding friendly port names and also quickly shows you which ports do not have friendly name assignments. (**show name** data comes from the running-config file.)
- **show interface <port-number>:** Displays the friendly port name, if any, along with the traffic statistics for that port. (The friendly port name data comes from the running-config file.)
- **show config:** Includes friendly port names in the per-port data of the resulting configuration listing. (**show config** data comes from the startup-config file.)

To List All Ports or Selected Ports with Their Friendly Port Names.

This command lists names assigned to a specific port.

Syntax: show name [port-list]

*Lists the friendly port name with its corresponding port number and port type. The **show name** command without a port list shows this data for all ports on the switch.*

For example:

```
ProCurve(config)# show name
Port Names
Port Type      Name
-----
A1  10/100TX    not assigned
A2  10/100TX    not assigned
A3  10/100TX    Bill_Smith@10.25.101.73
A4  10/100TX    not assigned
A5  10/100TX    Draft-Server:Trunk
A6  10/100TX    Draft-Server:Trunk
A7  10/100TX    Draft-Server:Trunk
A8  10/100TX    Draft-Server:Trunk
A9  10/100TX    not assigned
A10 10/100TX    not assigned
A11 10/100TX    not assigned
A12 10/100TX    not assigned
.    .
.    .
.    .
```

Ports Without "Friendly"

Friendly port names assigned in previous examples.

Figure 10-14. Example of Friendly Port Name Data for All Ports on the Switch

```
ProCurve(config)# show name A2,A3,A5
```

Port Names

Port : A2	Port Without a "Friendly" Name
Type : 10/100TX	
Name : not assigned	
Port : A3	
Type : 10/100TX	
Name : Bill_Smith@10.25.101.73	Friendly port names assigned in previous examples.
Port : A5	
Type : 10/100TX	
Name : Draft-Server:Trunk	

Figure 10-15. Example of Friendly Port Name Data for Specific Ports on the Switch

Including Friendly Port Names in Per-Port Statistics Listings. A friendly port name configured to a port is automatically included when you display the port's statistics output.

Syntax: show interface < port-number >
Includes the friendly port name with the port's traffic statistics listing.

For example, if you configure port A1 with the name "O'Connor_10.25.101.43", the show interface output for this port appears similar to the following:

```
ProCurve(config)# show interface A1
```

Status and Counters - Port Counters for port A1

Name : O'Connor@10.25.101.43	Friendly Port Name
Link Status : Up	
Bytes Rx : 894,568	Bytes Tx : 2470
Unicast Rx : 1179	Unicast Tx : 13
Bcast/Mcast Rx : 5280	Bcast/Mcast Tx : 13
FCS Rx : 36	Drops Tx : 0
Alignment Rx : 2	Collisions Tx : 0
Runts Rx : 0	Late Colln Tx : 0
Giants Rx : 0	Excessive Colln : 0
Total Rx Errors : 38	Deferred Tx : 0

Figure 10-16. Example of a Friendly Port Name in a Per-Port Statistics Listing

For a given port, if a friendly port name does not exist in the running-config file, the Name line in the above command output appears as:

Name : not assigned

To Search the Configuration for Ports with Friendly Port Names.

This option tells you which friendly port names have been saved to the startup-config file. (**show config** does not include ports that have only default settings in the startup-config file.)

Syntax: show config

Includes friendly port names in a listing of all interfaces (ports) configured with non-default settings. Excludes ports that have neither a friendly port name nor any other non-default configuration settings.

For example, if you configure port A1 with a friendly port name:

```
ProCurve(config)# int A1 name Print_Server@10.25.101.43
ProCurve(config)# write mem
ProCurve(config)# int A2 name Herbert's PC
ProCurve(config)# show config

Startup configuration:
; J4850A Configuration Editor; Created on release #E.08.30
hostname "HPswitch"
time daylight-time-rule None
no cdp run
interface A1
  name "Print_Server@10.25.101.43"
exit
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit
no aaa port-access authenticator active
```

This command sequence saves the friendly port name for port A1 in the startup-config file. The name entered for port A2 is not saved because it was executed after write memory.

Listing includes friendly port name for port A1 only.

In this case, **show config** lists only port A1. Executing **write mem** after entering the name for port A2, and then executing **show config** again would result in a listing that includes both

Figure 10-17. Example Listing of the Startup-Config File with a Friendly Port Name Configured (and Saved)

— *This page is intentionally unused.* —

Power Over Ethernet (PoE) Operation for the Series 5300xl Switches

Contents

PoE Operation on the Series 5300xl Switches	11-2
Introduction	11-2
PoE Terminology	11-3
Overview of Operation	11-4
Related Publications	11-4
General PoE Operation	11-5
Configuration Options	11-5
PD Support	11-6
Power Priority Operation	11-8
Configuring PoE Operation	11-10
Changing the PoE Port Priority Level	11-10
Disabling or Re-Enabling PoE Port Operation	11-11
Changing the Threshold for Generating a Power Notice	11-11
Configuring Optional PoE Port Identifiers	11-12
Viewing PoE Configuration and Status	11-15
Displaying the Switch's Global PoE Power Status	11-15
Displaying an Overview of PoE Status on All Ports	11-16
Displaying the PoE Status on Specific Ports	11-17
Planning and Implementing a PoE Configuration	11-19
Assigning PoE Ports to VLANs	11-19
Applying Security Features to PoE Configurations	11-20
Assigning Priority Policies to PoE Traffic	11-20
Calculating the Maximum Load for an xl PoE Module	11-21
PoE Operating Notes	11-23
PoE Event Log Messages	11-23
"Informational" PoE Event-Log Messages	11-23
"Warning" PoE Event-Log Messages	11-25

PoE Operation on the Series 5300xl Switches

The Power Over Ethernet (PoE) features described in this chapter operate on ProCurve Switch Series 5300xl devices running software release E.08.20 (or greater), with one or more ProCurve Switch xl PoE (J8161A) modules installed. Each PoE module must be connected to either of the following:

- ProCurve 600 Redundant and External Power (J8168A) Supply (HP 600 RPS/EPS)
- ProCurve 610 External Power (J8169A) Supply (HP 610 EPS)

(The ProCurve 3400cl/6400cl switches do not include PoE operation.)

Introduction

PoE technology allows IP telephones, wireless LAN access points, and other appliances to receive power and transfer data over existing LAN cabling. (For more on this topic, refer to edition 2 or later of the *ProCurve xl Modules Installation Guide* shipped with your optional J8161A Switch xl PoE Module (beginning in April, 2004).

PoE Terminology

Term	Use in this Manual
active PoE port	A PoE-enabled port connected to a PD requesting power.
priority class	Refers to the type of power prioritization where an xl PoE module uses Low (the default), High , and Critical priority assignments to determine which groups of ports will receive power. Note that power priority rules apply on a per-module basis, and only if PoE provisioning on a given module becomes oversubscribed.
EPS	External Power Supply; for example, an HP 600 ProCurve RPS/EPS or a ProCurve 610 EPS. An EPS device provides power to provision PoE ports on a module. See also "RPS", below.
MPS	Maintenance Power Signature; the signal a PD sends to the switch to indicate that the PD is connected and requires power. Refer to figure 11-11-4 on page 11-18.
Over-Subscribe	The state of a J8161A xl PoE module where there are more PDs requesting PoE power than the module has power to accommodate.
PD	Powered Device. This is an IEEE 802.3af-compliant device that receives its power through a direct connection to a 10/100Base-TX PoE RJ-45 port in an xl PoE module. Examples of PDs include Voice-over-IP (VoIP) telephones, wireless access points, and remote video cameras.
port-number priority	Refers to the type of power prioritization where, within a priority class, an xl PoE module assigns the highest priority to the lowest-numbered port in the module, the second-highest priority to the second lowest-numbered port in the module, and so-on. Note that power priority rules apply only if PoE provisioning on the module becomes oversubscribed.
PoE	Power-Over-Ethernet; the method by which PDs receive power from an xl PoE module (in compliance with the IEEE 802.3AF standard).
PSE	Power-Sourcing Equipment. A PSE, such as a J8161A xl PoE module installed in a Series 5300xl switch, provides power to IEEE 802.3AF-compliant PDs directly connected to the ports on the module. The xl PoE module is an <i>endpoint</i> PSE.
RPS	Redundant Power Supply; for example, the non-EPS operation of an HP 600 RPS/EPS. An RPS device provides power to a switch if the switch's internal power supply fails. RPS power does not provision PoE ports in modules installed in the 5300xl switches. See also "EPS", above.
RPS/EPS	A device that delivers redundant power to run a switch and external power to support PoE operation on a switch.
xl PoE Module	Refers to a ProCurve Switch xl PoE Module (J8161A).

Overview of Operation

A J8161A xl PoE module is a PSE device that receives PoE power from an external EPS device and distributes this power to the PDs connected to the xl PoE module's RJ-45 ports. The xl PoE module receives either 204 watts or 408 watts from the EPS, depending on whether the EPS is supporting one or two PSE devices.

Note

You can connect either a PoE device (PD) or a non-PoE device to a port configured for PoE operation on a J8161A xl PoE module.

Regarding Cat-5 cabling for PoE, the 802.3af standard allows either the spare pin/wire pairs or the data pin/wire pairs for PoE power transmission. A PoE module installed in a series 5300xl device supplies PoE power over the data pin/wire pairs. For more on this topic, refer to the *PoE Planning and Implementation Guide* (p/n 5990-6045, Nov. 2003 or later) available on the ProCurve Networking web site. (See “Getting Documentation From the Web” on page 1-6.)

Using the commands described in this chapter, you can:

- Configure a non-default power threshold for SNMP and Event Log reporting of PoE consumption on either all PoE ports on the switch or on all PoE ports in one or more PoE modules.
- Specify the port priority you want each xl PoE module to use for provisioning PoE power in the event that a given module's PoE resources become oversubscribed.
- Enable or disable PoE operation on individual ports. (In the default configuration, and with software release E.08.20 or greater installed, each xl PoE module installed in the switch enables PoE power on all 10/100-TX ports in the module, subject to PoE priority if the PoE resources on a given PoE xl module are oversubscribed.)
- Monitor PoE status and performance per module.

Related Publications

This chapter introduces general PoE operation, PoE configuration and monitoring commands, and Event Log messages related to PoE operation on a ProCurve Switch Series 5300xl device with one or more PoE modules installed and supported by the necessary external power supplies. The following two manuals provide further information:

- For information on installing a ProCurve Switch xl PoE Module (J8161A), refer to the *ProCurve Switch xl Modules Installation Guide* provided with the module.
- To help you plan and implement a PoE system in your network, refer to edition 2 or later of the *PoE Planning and Implementation Guide*, which is available from either of the following sources:
 - The Documentation CD-ROM (version 3.6 or greater) shipped with your Switch Series 5300xl device after April, 2004.
 - The ProCurve Networking web site at **www.procurve.com**. (Click on **technical support**, then **Product manuals (all)**.)

The latest version of any ProCurve product guide is always on the ProCurve Networking web site. See to “Getting Documentation From the Web” on page 1-6.

General PoE Operation

Configuration Options

In the default configuration, all 10/100Base-TX ports on an xl PoE module installed in the switch are configured to support PoE operation. You can:

- Disable or re-enable per-port PoE operation on individual ports to help control power usage and avoid oversubscribing PoE resources.
- Configure per-port priority for allocating power in case a PoE module becomes oversubscribed and must drop power for some lower-priority ports to support the demand on other, higher-priority ports.
- Configure one of the following:
 - A global power threshold that applies to all modules on the switch. This setting acts as a trigger for sending a notice when the PoE power consumption on any xl PoE module installed in the switch crosses the configured global threshold level. (Crossing the threshold level in either direction—PoE power usage either increasing or decreasing—triggers the notice.) The default setting is 80%.
 - A per-slot power threshold that applies to an individual xl PoE module installed in the designated slot. This setting acts as a trigger for sending a notice when the module in the specified slot exceeds or goes below a specific level of PoE power consumption.

Note

The ports on a PoE module support standard networking links and PoE links. Thus, you can connect either a non-PoE device or a PD to a PoE-enabled port without reconfiguring the port.

PD Support

An xl PoE module must have a minimum of 15.4 watts of unused PoE power available when you connect an 802.3af-compliant PD, regardless of how much power the PD actually uses. Depending on the amount of power the EPS device delivers to a specific xl PoE module, there may or may not always be enough power available to connect and support 802.3af PoE operation on all 24 10/100-TX ports. For example, if an EPS device is supporting only one xl PoE module and no other PSEs, then there will be sufficient power available for all ports on the module. However, if the same EPS is supporting both an xl PoE module and another ProCurve PSE device then, depending on the power demand placed on the module by the PDs you connect, it is possible to oversubscribe the available PoE power on the module. In this case, one or more PDs connected to the module will not have power. That is:

- **Sufficient PoE Power Available:** When a new PD connects to an xl PoE module in the switch, and if the module has a minimum of 15.4 watts of unused PoE power available, the module supplies power to the port for that PD.
- **Insufficient PoE Power Available:** When a new PD connects to an xl PoE module, and if the module does not have a minimum of 15.4 watts of unused PoE power already available:
 - If the new PD connects to a port “X” having a *higher* PoE priority than another port “Y” that is already supporting another PD on the same module, then the module removes PoE power from port “Y” and delivers it to port “X”. In this case the PD on port “Y” loses power and the PD on port “X” receives power.
 - If the new PD connects to a port “X” having a *lower* priority than all other PoE ports currently providing power to PDs on the same module, then the module does not deliver PoE power to port “X” until one or more PDs using higher priority ports are removed.

Note that once a PD connects to a PoE port and begins operating, the port retains only enough PoE power to support the PD’s operation. Unneeded power becomes available for supporting other PD connections. Thus, while 15.4 watts must be available for an xl PoE module on the switch to begin supplying power to a port with a PD connected, 15.4 watts per port is not continually required if the connected PD requires less power. For example,

with 20 watts of PoE power remaining available on a module, you can connect one new PD without losing power to any currently connected PDs on that module. If that PD draws only 3 watts, then 17 watts remain available and you can connect at least one more PD to that module without interrupting power to any other PoE devices connected to the same module. If the next PD you connect draws 5 watts, then only 12 watts remain unused. With only 12 unused watts available, if you then connect yet another PD to a higher-priority PoE port, then the lowest-priority port on the module loses PoE power and remains unpowered until the module once again has 15.4 or more watts available. (For information on power priority, refer to “Power Priority Operation” on page 11-8.)

Disconnecting a PD from a PoE port causes the module to stop providing PoE power to that port and makes the power available to any other PoE ports on the module that have PDs connected and waiting for power. If the PD demand for power on a module becomes greater than the PoE power available on the module, then the module transfers power from its lower-priority ports to its higher-priority ports. (Ports not currently providing power to PDs are not affected.)

Note

15.4 watts of available power is required for an xl PoE module on a switch to begin delivering power to a port, such as when a newly connected PD is detected or when power is released from higher-priority ports. Depending on power demands, lower-priority ports on a module with high PoE power demand may occasionally lose power due to the demands of higher-priority ports on the same module. (Refer to “Power Priority Operation” on page 11-8.)

Table 11-1. xl PoE Module Maximum Power Allocations

Power-Sourcing Equipment (PSE) Load on the EPS	Power to PoE Module from External EPS ¹	PoE Power Available for the xl PoE (J8161A) Module
One xl PoE Module Only	408 Watts	Maximum (15.4 W) available to all ports on the module.
Two PSE Devices (Two xl PoE Modules, or one xl PoE Module and One ProCurve PoE Stackable Switch)	204 Watts	Depending on the power demand from the PDs, lower priority ports may not be provisioned. Refer to “Calculating the Maximum Load for an xl PoE Module” on page 11-21.
¹ If a ProCurve EPS device is supplying PoE power to two PSE devices, then both PSE devices receive 204 watts. If a ProCurve EPS device is delivering PoE power to only one PSE device, then that device receives 408 watts.		

Power Priority Operation

When Does an xl PoE Module Prioritize Power Allocations? If an xl PoE module can provide power for all connected PD demand, it does not use its power priority settings to allocate power. However, if the PD power demand oversubscribes the available power, then the module prioritizes the power allocation to the ports that present a PD power demand. This causes the module to remove power from one or more lower-priority ports to meet the power demand on other, higher-priority ports. (This operation occurs, regardless of the order in which PDs connect to the module's PoE-enabled ports.) Note that each PoE xl module is a stand-alone priority domain. The switch does not prioritize one PoE module over another.

How Does an xl PoE Module Prioritize Power Allocations? xl PoE modules apply the following priority scheme:

- Using a *priority class* method, the module assigns a power priority of **Low** (the default), **High**, or **Critical** to each enabled PoE port.
- Using a *port-number priority* method, the module gives a lower-numbered port priority over a higher-numbered port within the same configured priority class.

Suppose, for example, that you configure the PoE priority for a module in slot C as shown in table 11-2.

Table 11-2. Example of PoE Priority Operation on an xl PoE Module

Port	Priority Setting	Configuration Command ¹ and Resulting Operation with PDs connected to Ports C3 Through C24
C3 - C17	Critical	<p>In this example, the following CLI command sets ports C3-C17 to Critical:</p> <pre>ProCurve(config)# interface c3-c17 power critical</pre> <p>The Critical priority class always receives power. If there is not enough power to provision PDs on all of the ports configured for this class, then no power goes to ports configured for High and Low priority. If there is enough power to provision PDs on only some of the critical-priority ports, then power is allocated to these ports in ascending order, beginning with the lowest-numbered port in the class, which, in this case, is port 3.</p>
C18 - C21	High	<p>In this example, the following CLI command sets ports C19-C22 to High:</p> <pre>ProCurve(config)# interface c19-c22 power high</pre> <p>The High priority class receives power only if all PDs on ports with a Critical priority setting are receiving power. If there is not enough power to provision PDs on all ports with a high priority, then no power goes to ports with a low priority. If there is enough power to provision PDs on only some of the high-priority ports, then power is allocated to these ports in ascending order, beginning, in this example, with port 18, until all available power is in use.</p>
C22 - C24	Low	<p>In this example, the CLI command sets ports C23-C24 to Low²:</p> <pre>ProCurve(config)# interface c23-c24 power low</pre> <p>This priority class receives power only if all PDs on ports with High and Critical priority settings are receiving power. If there is enough power to provision PDs on only some low-priority ports, then power is allocated to the ports in ascending order, beginning with the lowest-numbered port in the class (port 22, in this case), until all available power is in use.</p>
C1 - C2	- n/a -	<p>In this example, the CLI command disables PoE power on ports C1-C2:</p> <pre>ProCurve(config)# no interface c1-c2 power</pre> <p>There is no priority setting for the ports in this example.</p>
<p>¹ For a listing of PoE configuration commands, with descriptions, refer to “Configuring PoE Operation” on page 11-10.</p> <p>² In the default PoE configuration, the ports are already set to the low priority. In this case, the command is not necessary.</p>		

Configuring PoE Operation

In the default configuration, PoE support is enabled on the 10/100Base-TX ports in an xl PoE (J8161A) module installed on the switch. The default priority for all ports is **Low** and the default power notification threshold is **80** (%). Using the CLI, you can:

- Change the PoE priority level on individual PoE ports
- Disable or re-enable PoE operation on individual PoE ports
- Change the threshold for generating a power level notice

Changing the PoE Port Priority Level

Syntax: interface < port-list > power [critical | high | low]

Reconfigures the PoE priority level on < port-list >. For a given level, the module automatically prioritizes ports by port number (in ascending order). If there is not enough power available to provision all active PoE ports at a given priority level, then the lowest-numbered port at that level will be provisioned on a specific module first, and so on. An xl PoE module invokes configured PoE priorities only when it cannot provision all active PoE ports on that module.

- **Critical:** Specifies the highest-priority PoE support for < port-list >. The module provisions active PoE ports at this level before provisioning PoE ports at any other level.
- **High:** Specifies the second priority PoE support for < port-list >. The module provisions active PoE ports at this level before provisioning Low priority PoE ports.
- **Low** (the default): Specifies the third priority PoE support for < port-list >. The module provisions active PoE ports at this level only if there is power available after provisioning any active PoE ports at the higher priority levels.

*You can use one command to set the same priority level on PoE ports in multiple modules. For example, to configure the priority to **High** for ports c5-c10, C23-C24, D1-D10, and D12, you could use this command:*

```
ProCurve(config)# interface c5-c10,c23-c24,d1-d10,d12
```

Disabling or Re-Enabling PoE Port Operation

Syntax: [no] interface [e] < port-list > power

*Re-enables PoE operation on < port-list > and restores the priority setting in effect when PoE was disabled on < port-list >. The [no] form of the command disables PoE operation on < port-list >. (Default: All xl PoE ports on the module are enabled for PoE operation at **Low** priority.)*

Changing the Threshold for Generating a Power Notice

Syntax: power [slot < slot-identifier >] threshold < 1 - 99 >

This command specifies the PoE usage level (as a percentage of the PoE power available on a module) at which the switch generates a power usage notice. This notice appears as an SNMP trap and a corresponding Event Log message, and occurs when an xl PoE module's power consumption crosses the configured threshold value. That is, the switch generates a notice whenever the power consumption on a module either exceeds or drops below the specified percentage of the total PoE power available on the module.

This command configures the notification threshold for PoE power usage on either a global or per-module (slot) basis.

Without the [slot < slot-identifier >] option, the switch applies one power threshold setting on all PoE modules installed in the switch. For example, suppose slots A, B, and C each have an xl PoE module installed. In this case, executing the following command sets the global notification threshold to 70% of available PoE power.

```
ProCurve(config)# power threshold 70
```

With this setting, if an increasing PoE power demand crosses this threshold on the module in slot B, the switch sends an SNMP trap and generates this Event Log message:

*Slot B POE usage has exceeded threshold of 70 %.
If the switch is configured for debug logging, it also sends the Event Log message to the configured debug destination(s).*

On any PoE module, if an increasing PoE power load (1) exceeds the configured power threshold (which triggers the log message and SNMP trap), and then (2) later decreases and drops below the threshold again, the switch generates another SNMP trap, plus a message to the Event Log and any configured Debug destinations.

— Continued —

Syntax: power [slot < slot-identifier >] threshold < 1 - 99 > **(Continued)**

To continue the preceding example, if the PoE power usage on the xl PoE module in slot B drops below 70%, another SNMP trap is generated and you will see this message in the Event Log:

*Slot B POE usage is below threshold of 70 %.
For a message listing, refer to “PoE Event Log Messages” on page 11-23. (Default Global PoE Power Threshold: 80) By using the [slot < slot-identifier >] option, you can specify different notification thresholds for different xl PoE modules installed in the switch. For example, you could set the power threshold for a PoE module in slot “A” to 75% and the threshold for the module in slot “B” to 68% by executing the following two commands:*

```
ProCurve(config)# power slot a threshold 75  
ProCurve(config)# power slot b threshold 68
```

*Note that the last **threshold** command affecting a given slot supersedes the previous threshold command affecting the same slot. Thus, executing the following two commands in the order shown sets the threshold for the PoE module in slot “D” to 75%, but leaves the thresholds for any PoE modules in the other slots at 90%.*

```
ProCurve(config)# power threshold 90  
ProCurve(config)# power slot d threshold 75
```

(If you reverse the order of the above two commands, all PoE modules in the switch will have a threshold of 90%.)

Configuring Optional PoE Port Identifiers

The **Configured Type** field enables you to configure a unique identifier for PoE ports that helps to identify the intended use for a given PoE port. Such identifiers are useful when viewing PoE status with the following commands:

show power-management brief (page 11-16)

show power-management < port-list > (page 11-17)

To configure a unique identifier for one or more PoE ports, use the switch's **setmib** command to change the identifier setting in the switch's MIB (Management Information Base), as described in the following steps.

1. Use the **walkmib pethPsePortType.< slot-# >** command to determine the MIB-based port number for the port to which you want to assign a **Configured Type** identifier. On the 5300xl switches the slot numbering is as follows:

Slot	Slot Number Used in the MIB
A	1
B	2
C	3
D	4
E*	5
F*	6
G*	7
H*	8

*5308xl only.

Note that in the MIB, 26 port numbers are assigned to each slot designation. Thus, for example, with PoE modules in slots “A” and “B”, the actual, corresponding port numbers will be 1-24 and 27-50, respectively. (The port numbers “25”, “26”, “51”, and “52” are reserved.)

2. Use the **setmib pethPsePortType.< slot-# >.< port-# > -D < identifier-string >** command to configure the identifier you want for a specific port.

For example, suppose that you have a PoE xl Module installed in slot B and want to assign the identifier “Wireless-1” to port 1 in this slot. To do so, you would use the following commands:

Power Over Ethernet (PoE) Operation for the Series 5300xl Switches

Configuring PoE Operation

```

ProCurve(config)# walkmib pethPsePortType.2
pethPsePortType.2.27 =
pethPsePortType.2.28 =
pethPsePortType.2.29 =
pethPsePortType.2.30 =
.      .      .
.      .      .

ProCurve(config)# setmib pethPsePortType.2.27 -D Wireless-1
pethPsePortType.2.27 = Wireless-1

ProCurve(config)# show power-management brief

Status and Counters - Port Power Status

Port | Power Enable | Priority | Configured Type | Detection Status | Power Class
-----+-----
B1   | Yes          | Low     | Wireless-1      | Searching         | 0
B2   | Yes          | Low     |                  | Searching         | 0
B3   | Yes          | Low     |                  | Searching         | 0
B4   | Yes          | Low     |                  | Searching         | 0
.    | .            | .       |                  | .                | .
.    | .            | .       |                  | .                | .
.    | .            | .       |                  | .                | .

```

Lists port numbers used by the MIB for slot "B".

MIB Designation for Port B1

Command to configure "Wireless-1" as the **Configured Type** identifier for port B1.

CLI response indicates successful command execution.

"Show" command lists the new **Configured Type** identifier.

Figure 11-1.Example of using the MIB To Configure a "Configured Type" Identifier for a Port

To remove a Configured Type identifier, use the setmib command with a blank space enclosed in quotes. For example, to return port B2 in the above figure to a null setting, use this command:

```
ProCurve(config)# setmib pethPsePortType.2.27 -D " "
```

For more on displaying PoE configuration and status, refer to "Viewing PoE Configuration and Status" on page 11-15.

Viewing PoE Configuration and Status

Displaying the Switch's Global PoE Power Status

Syntax: show power-management

Displays the switch's global PoE power status, including:

- **Maximum Power:** *Lists the maximum PoE wattage available to provision active PoE ports on the switch.*
- **Power In Use:** *Lists the amount of PoE power presently in use.*
- **Operational Status:** *Indicates whether PoE power is available on the switch. (Default: **On** ; shows **Off** if PoE power is not available. Shows **Faulty** if internal or external PoE power is oversubscribed or faulty.)*
- **Usage Threshold (%):** *Lists the configured percentage of available PoE power provisioning the switch must exceed to generate a usage notice in the form of an Event Log message and an SNMP trap. If this event is followed by a drop in power provisioning below the threshold, the switch generates another SNMP trap and Event Log message. Event Log messages are also sent to any optionally configured debug destinations. (Default: **80%**)*

For example, in the default PoE configuration, when the switch is running with several ports on the xl PoE modules in slots C and D supporting PD loads, **show power-management** displays data similar to the following:

```
ProCurve(config)# show power-management

Status and Counters - System Power Status for slot C

Maximum Power   : 204 W           Operational Status : On
Power In Use    : 46 W +/- 6 W    Usage Threshold (%) : 80

Status and Counters - System Power Status for slot D

Maximum Power   : 204 W           Operational Status : On
Power In Use    : 34 W +/- 6 W    Usage Threshold (%) : 80
```

Figure 11-2. Example of Show Power-Management Output

Displaying an Overview of PoE Status on All Ports

Syntax: show power-management brief

Displays the following port power status:

- **Port:** Lists all PoE-capable ports on the switch.
- **Power Enable:** Shows **Yes** for ports enabled to support PoE (the default) and **No** for ports on which PoE is disabled.
- **Priority:** Lists the power priority (**Low**, **High**, and **Critical**) configured on ports enabled for PoE. (For more on this topic, refer to the power command description under “Configuring PoE Operation” on page 11-10.)
- **Configured Type:** If configured, shows the user-specified identifier for the port. If not configured, the field is empty. Refer to “Configuring Optional PoE Port Identifiers” on page 11-12.
- **Detection Status:**
 - **Searching:** The port is trying to detect a PD connection.
 - **Delivering:** The port is delivering power to a PD.
 - **Disabled:** On the indicated port, either PoE support is disabled or PoE power is enabled but the xl PoE module does not have enough power available to supply the port’s power needs.
 - **Fault:** The switch detects a problem with the connected PD.
- **Power Class:** Shows the 802.3af power class of the PD detected on the indicated port. Classes include:

0: 0.44w to 12.95w	3: 6.49w to 12.95w
1: 0.44w to 3.84w	4: reserved
2: 3.84w to 6.49w	

For example, **show power-management brief** displays this output:

```
ProCurve(config)# show power-management brief
```

Status and Counters - Port Power Status

Port	Power Enable	Priority	Configured Type	Detection Status	Power Class
C1	Yes	Critical	Telephone	Delivering	1
C2	Yes	Critical	Telephone	Delivering	1
C3	Yes	High	Wireless	Delivering	3
C4	Yes	High	Wireless	Delivering	3
C5	Yes	Low		Searching	0
C6	Yes	Low		Searching	0
C7	Yes	Low		Searching	0
C8	Yes	Low		Searching	0
.
.

Ports C1 through C4 are delivering power. The remaining ports are available to supply power, but currently do not detect a connected PD.

Figure 11-3. Example of Show Power-Management Brief Output

Displaying the PoE Status on Specific Ports

Syntax: show power-management < port-list >

Displays the following PoE status and statistics (since the last reboot) for each port in < port-list >:

- **Power Enable:** Shows **Yes** for ports enabled to support PoE (the default) and **No** for ports on which PoE is disabled. Note that for ports on which power is disabled, this is the only field displayed by **show power-management < port-list >**.
- **Priority:** Lists the power priority (**Low**, **High**, and **Critical**) configured on ports enabled for PoE. (For more on this topic, refer to the power command description under “Configuring PoE Operation” on page 11-10.)
- **Detection Status:**
 - **Searching:** The port is available to support a PD.
 - **Delivering:** The port is delivering power to a PD.
 - **Disabled:** PoE power is enabled on the port but the xl PoE module does not have enough power available to supply the port’s power needs.
- **Fault:** The switch detects a problem with the connected PD.
- **Over Current Cnt:** Shows the number of times a connected PD has attempted to draw more than 15.4 watts. Each occurrence generates an Event Log message.

— Continued —

Syntax: show power-management < port-list> *(Continued)*

- **Power Denied Cnt:** Shows the number of times PDs requesting power on the port have been denied due to insufficient power available. Each occurrence generates an Event Log message.
- **Voltage:** The total voltage, in dV, being delivered to PDs.
- **Power:** The total power, in mW, being delivered to PDs.
- **Configured Type:** If configured, shows the user-specified identifier for the port. If not configured, the field is empty. Refer to “Configuring Optional PoE Port Identifiers” on page 11-12.
- **Power Class:** Shows the power class of the PD detected on the indicated port. Classes include:

0: 0.44w to 12.95w	2: 3.84w to 6.49w	4: reserved
1: 0.44w to 3.84w	3: 6.49w to 12.95w	
- **MPS Absent Cnt:** This value shows the number of times a detected PD has no longer requested power from the port. Each occurrence generates an Event Log message. (“MPS” refers to the “Maintenance Power Signature”. Refer to “PoE Terminology” on page 11-3.)
- **Short Cnt:** Shows the number of times the switch provided insufficient current to a connected PD.
- **Current:** The total current, in mA, being delivered to PDs.

For example, if you wanted to view the PoE status of ports C1 and D5, you would use **show power-management c1,d5** to display the data:

```
ProCurve(config)# show power-management d4-d5
```

Status and Counters - Port Power Status for port D4

Power Enable : Yes	
Priority : Low	Configured Type :
Detection Status : Delivering	Power Class : 0
Over Current Cnt : 0	MPS Absent Cnt : 0
Power Denied Cnt : 0	Short Cnt : 0
Voltage : 492 dV	Current : 52 mA
Power : 14210 mW	

Status and Counters - Port Power Status for port D5

Power Enable : No	
-------------------	--

Example of command output for a port on which power is enabled.

Example of command output for a port on which power is disabled.

Figure 11-4. Example of Show Power-Management < port-list> Output

Planning and Implementing a PoE Configuration

This section provides an overview of some considerations for planning a PoE application. For additional information on this topic, refer to the *ProCurve PoE Planning and Implementation Guide*.

Some of the elements you may want to consider for a PoE installation include:

- Port assignments to VLANs
- Use of security features
- Power requirements

This section can help you to plan your PoE installation. If you use multiple VLANs in your network, or if you have concerns about network security, you should read the first two topics. If your PoE installation comes close to (or is likely to exceed) the system's ability to supply power to all devices that may request it, then you should also read the third topic. (If it is unlikely that your installation will even approach a full utilization of the PoE power available, then you may find it unnecessary to spend much time on calculating PoE power scenarios.)

Assigning PoE Ports to VLANs

If your network includes VLANs, you may want to assign various PoE-configured ports to specific VLANs. For example, if you are using PoE telephones in your network, you may want to assign ports used for telephone access to a VLAN reserved for telephone traffic.

Applying Security Features to PoE Configurations

You can utilize security features built into the switch to control device or user access to the network through PoE ports in the same way as non-PoE ports.

- **MAC Address Security:** Using Port Security, you can configure each switch port with a unique list of MAC addresses for devices that are authorized to access the network through that port. For more information, refer to the chapter titled "Configuring and Monitoring Port Security" in the *Access Security Guide* for your switch.
- **Username/Password Security:** If you are connecting a device that allows you to enter a username and password that is forwarded to a networked server for authentication, then you can also configure the following security features:
 - Local username and password
 - TACACS+
 - RADIUS Authentication and Accounting
 - 802.1X Authentication

For more information on security options, refer to the latest edition of the *Access Security Guide* for your switch. (The ProCurve Networking web site offers the latest version of all ProCurve product publications. Refer to "Getting Documentation from the Web" in chapter 1, "Getting Started".)

Assigning Priority Policies to PoE Traffic

You can use the configurable QoS (Quality of Service) features in the switch to create prioritization policies for traffic moving through PoE ports. Table 11-3 lists the available classifiers and their order of precedence.

Table 11-3. Classifiers for Prioritizing Outbound Packets

Priority	QoS Classifier
1	UDP/TCP Application Type (port)
2	Device Priority (destination or source IP address)
3	IP Type of Service (ToS) field (IP packets only)
4	Protocol Priority (IP, IPX, ARP, DEC LAT, AppleTalk, SNA, and NetBeui)
5	VLAN Priority
6	Incoming source-port on the switch
7	Incoming 802.1p priority (present in tagged VLAN environments)

For more on this topic, refer to the chapter titled “Quality of Service: Managing Bandwidth More Effectively” in the *Advanced Traffic Management Guide* for your switch.

Calculating the Maximum Load for an xl PoE Module

Since the full PoE load on an xl PoE module receiving 408 watts (from an EPS supporting only that module) cannot exceed 369.6 watts (24 ports with a maximum of 15.4 watts per port), there is no concern for overloading the module’s PoE capacity. However, for xl PoE modules receiving 204 watts due to EPS power-sharing with another PoE device, it is possible to exceed the maximum supportable load. Also, when planning the PoE load, the following factors apply per-module:

- When first connecting an appliance to a PoE port, the xl PoE module must have a minimum of 15.4 watts of available PoE power. PoE power is “available” if it is either not currently in use or can be acquired by (automatically) removing PoE power from another, lower-priority port on the module. (See to “PD Support” on page 11-6.)
- After an appliance is connected to a PoE port, the switch reduces the power requirement for that port from the initial 15.4 watts to the actual power level the appliance requires.

Thus, after you have connected all but the last planned appliance to a PoE module, there must be a minimum of 15.4 watts of unused PoE power available on the module to support adding the final appliance. That is, where:

n = the total number of appliances you want to connect to one xl PoE module

and

w = the total PoE power required to operate $(n - 1)$ appliances

then, the following must be true:

$$w + 15.4 \leq 204$$

or

$$(204 - 15.4) \geq w$$

For example, suppose you have 24 identical appliances to connect to an xl PoE module receiving 204 watts of PoE power. For this example, each appliance requires 8.3 watts to operate. In this case, the module would support only 23 of these appliances at any given time because there would not be

enough unused power to meet the minimum of 15.4 watts required to support the initial bring-up of the 24th appliance. That is, $204 - (23 \times 8.3) = 13.1$. Because the module provisions power on the basis of the priority scheme described on page 11-10 (under the **interface < port-list > power [critical | high | low]** syntax), you can still fully populate the module with appliances. In this case, the lowest-priority port will not receive power unless an appliance in a higher-priority port is disconnected.

There is also a scenario where a device on a lower-priority port can experience a power cycle (temporarily lose power) while a higher-priority port is bringing up a PoE device. Suppose, for example, that:

1. An xl PoE module in slot B, with all ports configured at the default **Low** priority, is receiving 204 watts of power from an EPS.
2. The 21 PoE devices on ports B2 - B22 draw 9.0 watts of power each ($9 \times 21 = 189$), leaving 15.0 watts unused, which is less than the 15.4 watt minimum needed to add another PoE device to the module. (Refer to “PD Support” on page 11-6.)
3. The system operator plugs a 7.0-watt PoE device into port B1, which is the highest-priority port in slot B.

In the above scenario, there is less than 15.4 watts available to support the initial bringup of the newly installed device on port B1. As a result, port B22 (the lowest-priority port on the module) temporarily loses power so that there is enough power to add the new device on port B1. After the new device begins operation, the power demand on port B1 drops to 7 watts. At this point, there are 20 devices consuming 9 watts each, and 1 device consuming 7 watts, for a total of 187 watts, and the module now has 17 watts of unused power available. Since this exceeds the minimum of 15.4 watts required to bring up any PoE device, there is now enough power available to bring back up the device on port B22.

PoE Operating Notes

- Simply disabling a PoE port does *not* affect power delivery through that port. To cycle the power on a PD receiving power from a PoE port on the switch, disable, then re-enable the power to that port. For example, to cycle the power on a PoE device connected to port 1 on an xl PoE module installed in slot D:

```
ProCurve(config)# no interface d1 power
ProCurve(config)# interface d1 power
```

PoE Event Log Messages

PoE operation generates these Event Log messages. You can also configure the switch to send these messages to a configured debug destination (terminal device or SyslogD server).

“Informational” PoE Event-Log Messages

I <MM/DD/YY> <HH:MM:SS> <chassis|ports>:

Message header, with severity, date, system time, and system module type (chassis or ports). For more information on Event Log operation, including severity indicators, refer to “Using the Event Log To Identify Problem Sources” on page C-27.

Slot <slot-id> Ext Power Supply connected, supplying
<actual-power> W of <avail-power> W max.

The switch detected an EPS (External Power Supply) on the indicated slot and began receiving the wattage indicated by <actual-power>. The <avail-power> field indicates the maximum power (wattage) the detected EPS is capable of delivering.

Slot <slot-id> Ext Power Supply disconnected

The indicated slot has lost contact with an external power supply.

Slot <*slot-id*> POE usage is below configured threshold
of <1 - 99> %

*Indicates that POE usage on the module in the indicated slot has decreased below the threshold specified by the last execution of the **powerthreshold** command affecting that module. This message occurs if, after the last reboot, the PoE demand on the module exceeded the power threshold and then later dropped below the threshold value.*

port <*port-id*> applying power to PD.

A PoE device is connected to the indicated port and receiving power.

port <*port-id*> PD detected.

The switch has detected a PoE device connected to the indicated port.

“Warning” PoE Event-Log Messages

W <MM/DD/YY> <HH:MM:SS> chassis:

Message header, with severity, date, system time, and system module type. For more information on Event Log operation, including severity indicators, refer to “Using the Event Log To Identify Problem Sources” on page C-27.”

Slot <slot-id> Ext Power Supply connected but not responding.

The switch detects an external power supply on the module in the indicated slot, but is not receiving power from the device.

Slot <slot-id> Ext Power Supply failure: <fault-type>
Failures:

Indicates an external power supply failure for the module in the indicated slot, where <fault-type> is one of the following:

- *Over Current fault: The external power supply reported a fault condition. Contact your ProCurve support representative.*
- *Fan fault: A fan in an external power supply has failed.*
- *Temperature fault: The operating temperature in an external power supply has exceeded the normal operating range.*
- *50V fault: The external power supply reported a fault condition. Contact your ProCurve support representative.*
- *12V fault: The external power supply reported a fault condition. Contact your ProCurve support representative.*

Slot <slot-id> POE usage has exceeded threshold of
<1-99> %

*Indicates that POE usage in the indicated slot has exceeded the configured threshold for the module, as specified by the last execution of the **power threshold** or **power slot <slot-id> threshold** command. (Note that the switch also generates an SNMP trap for this event.)*

Port <port-id> PD Denied power due to insufficient power allocation.

There is insufficient power available to power the PD on the indicated port and the port does not have sufficient PoE priority to take power from another active PoE port.

Port <*port-id*> PD Invalid Signature indication.

The switch has detected a non-802.3af-compliant device on the indicated port. This message appears for all non-802.3af devices connected to the port, such as other switches, PC-NICs, etc.

Port <*port-id*> PD MPS Absent indication.

The switch no longer detects a device on <port-id>. The device may have been disconnected, powered down, or stopped functioning.

Port <*port-id*> PD Other Fault indication.

There is a problem with the PD connected to the port.

Port <*port-id*> PD Over Current indication.

The PD connected to <port-id> has requested more than 15.4 watts of power. This may indicate a short-circuit or other problem in the PD.

Access Controller xl Module for the Series 5300xl Switches

Contents

Introduction	12-2
General Operation	12-2
Related Publications	12-2
Terminology	12-3
Access Controller xl Module Overview	12-4
Module Operation	12-4
Using 5300xl Features with the Access Controller xl Module ...	12-6
Routing Infrastructure Support	12-9
Using 5300xl Switch Network Address Translation with the ACM	12-10
The Role of VLANs	12-11
Client VLANs	12-11
Static VLAN Features Supported on Client VLANs	12-12
General Operating Rules	12-13
Configuring the ACM on the Network	12-13
Configuring the Access Controller xl Module	12-15
Configuring Downlink Client Ports	12-15
Changing the VLAN-Base	12-17
Configuring Client VLANs	12-18
Configuring Uplink Network Ports	12-18
Configuring the Uplink VLAN	12-18
ACM Configuration Commands Summary and Syntax	12-19
Configuration Context Command Syntax	12-19
Access Controller Context Command Syntax	12-21
Displaying Access Controller xl Status from the 5300xl CLI ...	12-23
ACM Display Commands Summary and Syntax	12-23

Configuration Context Command Syntax	12-24
Access Controller Context Command Syntax	12-25
Managing the ACM	12-26
Using the ACM's Extended CLI	12-26
Downloading New Software to the Module	12-29
Resetting the Module to Factory Defaults	12-29
Operating Notes	12-30
BIOS POST Event Log Messages	12-31

Introduction

The ProCurve Access Controller xl Module (ACM) enables secure, mobile user access to appropriate network services on any ProCurve Series 5300xl switch. This modular addition to the 5300xl switch offers a unique approach to integrating identity-based user access control, wireless data privacy and secure roaming with the flexibility of a full-featured intelligent edge switch. Centrally configured and managed access policies provide identity-based access control to wired and wireless users.

Note

The 5300xl switch software must be updated to version E.09.21 or later. The Access Control Server 740wl or the Integrated Access Manager 760wl must use software version 4.1.3.93 or later.

General Operation

The Access Controller xl Module (J8162A) uses ports on a 5300xl switch to pass wired and wireless traffic to and from the network using authentication and rights administration policies from an Access Control Server 740wl or an Integrated Access Manager 760wl. Up to two ACMs may be used in a single 5300xl switch. Once the ACM is installed in the switch, connected to the Access Control Server (740wl or 760wl), and configured for operation, it is managed from the Administrative Console of the Secure Access 740wl or 760wl products.

Related Publications

This chapter introduces Access Controller xl Module operation, configuration, and monitoring. The following two manuals provide further information:

- For information on installing the ACM, refer to the *ProCurve xl Modules Installation Guide* provided with the module.
- To help you manage and configure the ACM in your network, refer to the *ProCurve Secure Access 700wl Series Management and Configuration Guide*, which is available from either of the following sources:
 - The Documentation CD-ROM shipped with your module

- The ProCurve Networking Web site at www.procurve.com. (Click on **Technical support**, then **Product manuals (all)**.)

Terminology

Term	Use in this Manual
Access Control Server	<p>A centralized resource on the network that provides services, such as authentication management, mobility management (roaming support), policy management, and system monitoring and reporting, to the connected Access Controllers.</p> <p>The Access Control Server is deployed as a dedicated control function and does not sit in the user data path. The Secure Access 700wl Series has two products that provide this capability: the ProCurve Access Control Server 740wl and the Integrated Access Manager 760wl.</p>
Client	A device looking to access the network.
Client VLAN	A special VLAN created to handle downlink client port traffic for the ACM. Includes the downlink client port (with untagged VLAN membership) and the downlink port (<slot-id>DP) (with tagged VLAN membership).
Downlink Client Ports	Series 5300xl switch ports assigned as an untagged member to a client VLAN to supply client connectivity.
Downlink Port	The internal port that carries client traffic to and from the ACM. This port is identified by the slot ID where the module is installed, combined with 'DP'.
Integrated Access Manager 760wl	Combines the functionality of the ProCurve Access Controller 720wl and the ProCurve Access Control Server 740wl in a single device.
Uplink Port	The internal port that carries ACM traffic to and from the network. Must be an untagged member of a non-client VLAN. This port is identified by the slot ID where the module is installed, combined with 'UP.' For example, CUP is the uplink port for an ACM installed in slot C of a 5300xl switch.
Uplink Network Ports	Any 5300xl port that is a member of the uplink VLAN.
Uplink VLAN	The VLAN containing the uplink port as an untagged member. By default, this is the DEFAULT_VLAN on the 5300xl switch.

Access Controller xl Module Overview

The Access Controller xl Module adds new wireless security and access control capabilities to the 5300xl switch. The module supplies identity-based user access control to specific network services, wireless data privacy with VPN services, and application persistence across subnet boundaries at the edge of the network, where users connect. Centrally managed from the ProCurve Secure Access Control Server 740wl or Integrated Access Manager 760wl, the Access Controller xl Module provides hassle-free access while maintaining a high level of security.

Module Operation

Figure 12-12-1 below presents the module's key components. Each component is then discussed.

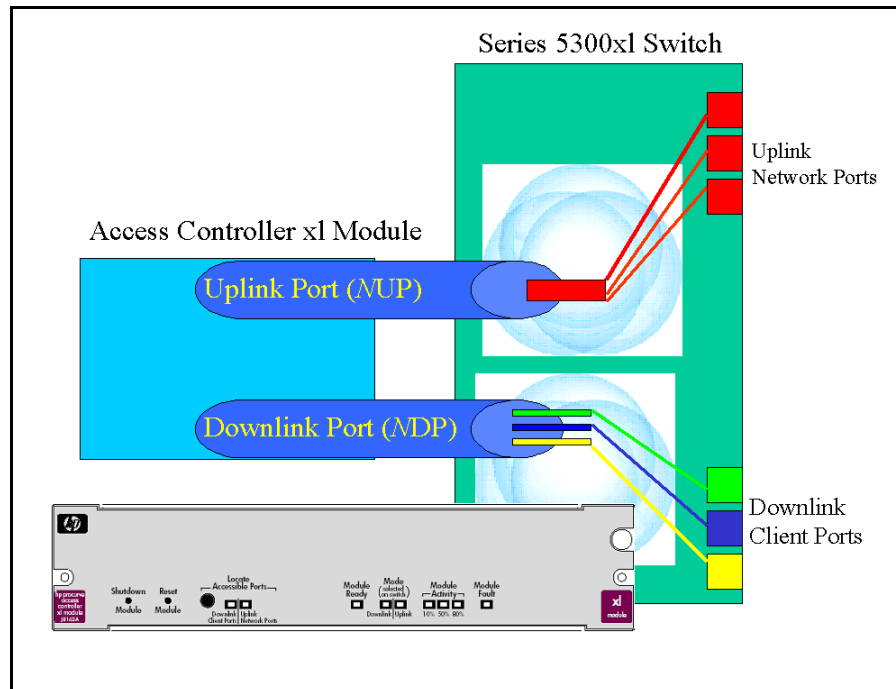


Figure 12-1. The Access Controller xl Module Conceptual View

The Access Controller xl Module has no external ports, as shown in Figure 12-12-1. The module uses ports on the 5300xl switch through two internal ports, the uplink port and the downlink port. Clients, typically connecting through an access point, connect to 5300xl ports defined as downlink client ports. The internal uplink port passes network traffic through other 5300xl ports, which are external uplink network ports. VLANs are used to direct traffic to and from the ACM.

For an explanation of the module's features and LEDs, see the *ProCurve xl Modules Installation Guide*.

Note

Uplink and downlink port names depend on the switch slot where the module is installed. When the module is in switch slot A, 'N' is 'A' in Figure 12-12-1. The uplink port for the module is AUP; the downlink port is ADP.

The following steps are required to add an ACM to your network:

1. Install an Access Control Server 740wl or Integrated Access Manager 760wl in the network, or identify an existing 740wl or 760wl to be used with the ACM.
2. Having identified the Access Control Server 740wl or Integrated Access Manager 760wl to be used with the ACM, note its IP address. To operate, the ACM must establish secure communications with the Access Control Server or Integrated Access Manager.

The shared secret configured on the 740wl/760wl's is also needed. If you are already using a 760wl, you may not have configured a shared secret. See "Editing the Access Control Server Configuration" in the *ProCurve Secure Access 700wl Series Management and Configuration Guide*, available on the Documentation CD-ROM shipped with your module or from the ProCurve Networking Web site at

www.procurve.com (Click on **Technical support**, then **Product manuals (all)**).

3. Install the ACM in a slot on the 5300xl switch. Once the Module Ready LED is on, the ACM requires an IP address. By default, the ACM uses DHCP. The IP address also can be set manually. The uplink port must be an untagged member of a VLAN that can communicate with the 740wl or 760wl. The ACM establishes communication with the 740wl/760wl, using the IP address and the shared secret from step 2 above. See the *ProCurve xl Modules Installation Guide* for details.

4. Configure downlink client ports, client VLANs, uplink network ports, and the uplink VLAN on the 5300xl switch. Configure access and user/group policy rights on the 740wl/760wl to support and manage clients and client traffic through the ACM.
5. Manage and monitor the ACM using the Administrative Console on a 740wl or 760wl.

There are specific installation and operational requirements for this device as a module in a Series 5300xl switch. The following sections describe how the module operates and how it is configured for use.

Using 5300xl Features with the Access Controller xl Module

As the ACM uses special ports and VLANs to provide access security to wireless devices, not all of the features of the 5300xl switch are applicable. For example, features that provide an alternative means of authentication are not supported on ACM downlink client ports.

Some 5300xl configurations are not allowed by the Command Line Interface (CLI). When a CLI command fails, a message is displayed explaining why. Warning messages are issued when an operation could potentially cause problems managing traffic through the ACM. For example, if a downlink client port is assigned to a non-client VLAN, traffic could enter the network without first being authenticated and assigned specific access rights by the ACM. In this case, a warning message is issued stating that the port is a member of a client VLAN. In some cases Log messages are also created when an operation is done, noting the potential conflict with ACM operation.

Note

5300xl switch ports that are not used by the Access Controller xl Module (that is, they are not downlink client ports, or members of client VLANs) continue to operate as regular 5300xl ports. Their operation is not affected.

The table below presents the 5300xl switch features that are not supported for use with an ACM module.

Feature	Uplink Port	Downlink Port	Downlink Client Ports	Client VLANs	Explanation
802.1X	x	x	x		Not allowed.
ACL				x	Has no effect if assigned. Warning issued
Configuring IP Addresses				x	Not allowed.
DHCP/DHCP Relay				x	Not allowed.
IP Helper Address				x	Not allowed.
Flow Control					Not supported across an ACM.
GVRP	x	x	x	x	GVRP cannot be enabled on an uplink, downlink, or downlink client port. A port in a GVRP VLAN cannot be added to a client VLAN. If GVRP is enabled on a port when it is added to a client VLAN, it is disabled.
IGMP		x	x	x	IGMP cannot be enabled on client VLANs. As a result, it cannot be enabled on downlink client ports.
Interface Monitoring (Port Mirroring)	x	x	x		Cannot be used as a monitoring port.
Interface Provisioning:					
Speed	x	x			Fixed at 1000Mbps.
Duplex	x	x			Fixed at Full-Duplex.
Flow-Control	x	x			Not allowed.
Auto-MDIX mode	x	x			Not allowed.
'x' indicates that the feature is not supported.					

Feature	Uplink Port	Downlink Port	Downlink Client Ports	Client VLANs	Explanation
IP Routing/ Multicast Routing				x	No routing is done.
				x	Not allowed.
IP Stacking					Not supported across an ACM.
IRDP				x	Not allowed.
Link Test	x	x			Test packets not supported across an ACM.
LLDP	x	x			Set to off.
MAC Auth	x	x	x		Not allowed.
Meshing	x	x	x		Not allowed
				x	Mesh ports cannot be a member of a client VLAN.
MSTP (802.1s)					An MSTP region may not span across an ACM.
OSPF				x	Not allowed.
PIM				x	Not allowed.
RIP				x	Not allowed.
Static VLANs					See table 12-9-1 below.
Trunking^a:					
LACP	x	x	x	x	Not allowed.
Virus Throttling		x	x	x	Not supported.
Web Auth	x	x	x		Not allowed.
XRRP				x	Not allowed.
'x' indicates that the feature is not supported.					

- a. A 5300xl switch trunk group that is configured using the **trunk** option, can be added to a **client VLAN**.

Routing Infrastructure Support

The ACM uses IP to communicate with Access Control Server 740wls, Integrated Access Managers 760wls and Access Controller 720wls. The default gateway must be set up correctly if there is a router in the communications path. Figure 12-12-2 shows an ACM communicating with its 740wl/760wl through a router.

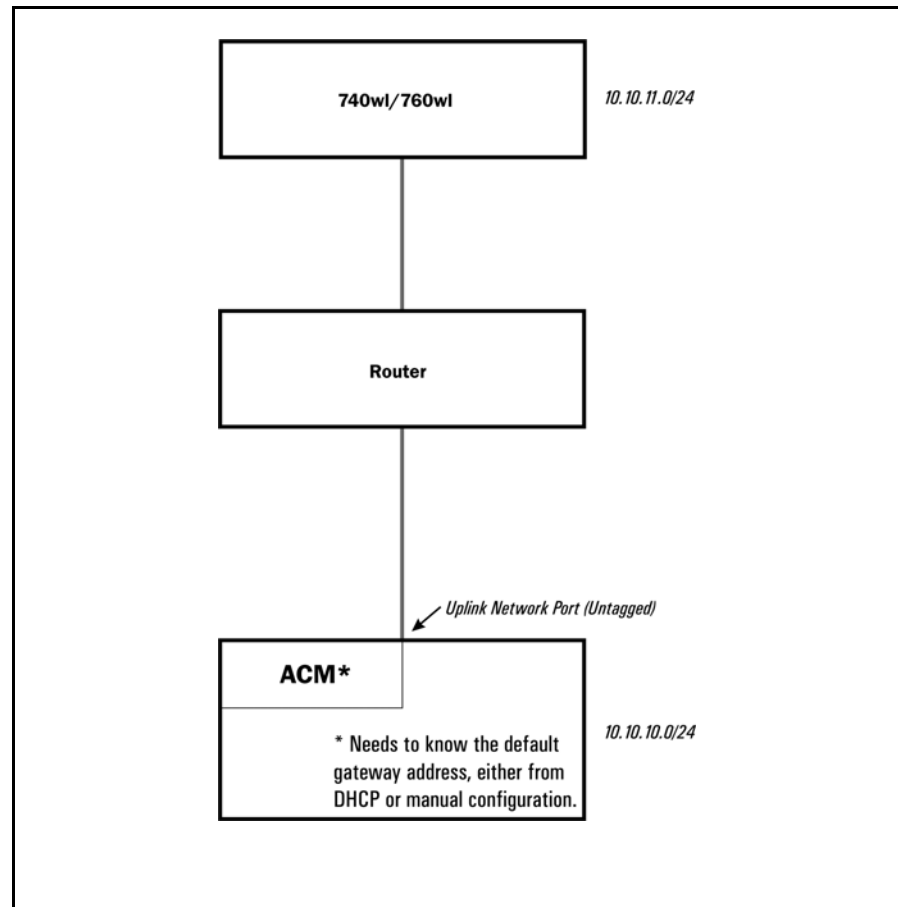


Figure 12-2. Accessing a 740wl or 760wl Through a Router

The ACM does not support any routing infrastructure attached to a downlink client port. Figure 12-12-3 below shows how an ACM can be used to communicate with a lower-level, non-routed network structure through a downlink client port.

Access Controller xl Module for the Series 5300xl Switches
Using 5300xl Features with the Access Controller xl Module

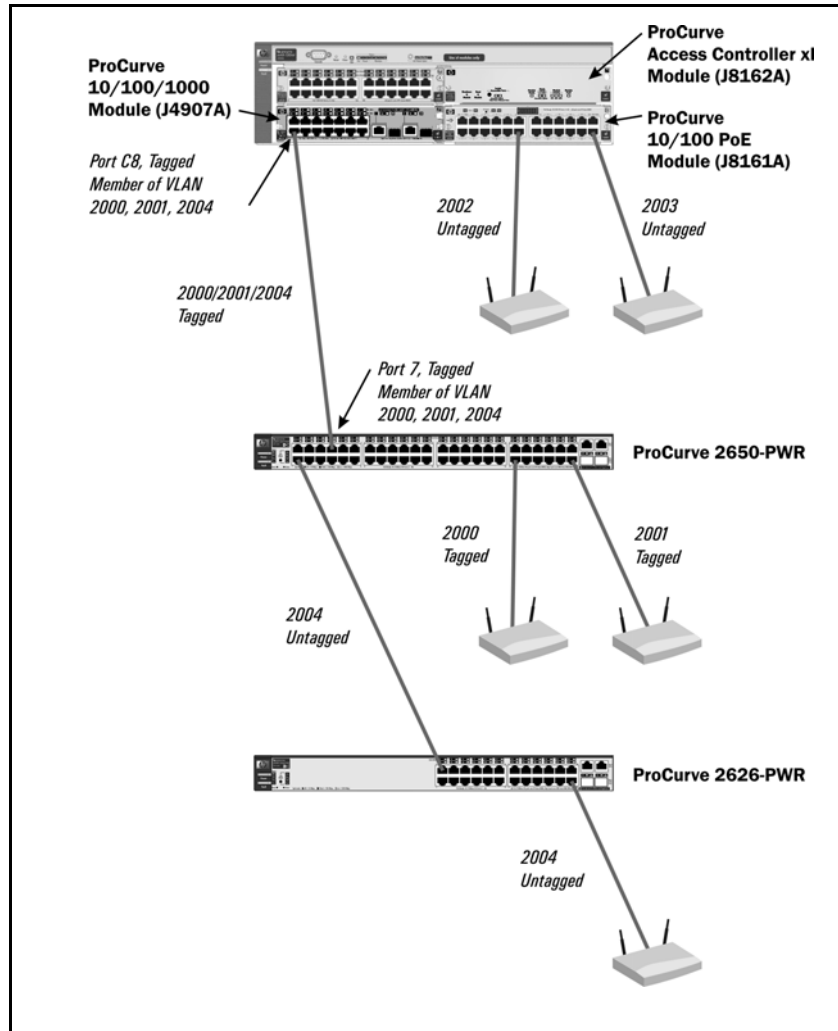


Figure 12-3. A Downlink Client Port with a Non-Routed Network Structure

Using 5300xl Switch Network Address Translation with the ACM

The Secure Access 700wl series products and the ACM provide network address translation for client traffic. The 5300xl switch's network address translation feature is not recommended for use with the ACM.

The Role of VLANs

VLANs are used by the Access Controller xl Module to manage client traffic through the switch. Downlink client ports, connecting to access points, either directly or through an intermediate network, are assigned as untagged members to a unique VLAN that also includes the downlink port as a tagged member. Traffic from the downlink client port, passing inbound through the downlink port on its way to the Access Controller xl Module, is normally tagged with the downlink client port's VLAN ID (VID), except when traffic is being bridged (see “Operating Notes” on page 12-31). The correct authentication policies and access policies are applied to this inbound client traffic by the Access Controller xl Module, based, in part, on the VLAN tag the traffic carries.

In a similar fashion, ACM traffic outbound to the network uses a VLAN to connect to the correct switch port. The uplink network port is an untagged member of the uplink VLAN, which by default is the 5300xl **DEFAULT_VLAN**. All switch ports that belong to the uplink VLAN are uplink network ports. The uplink VLAN may be changed by creating a new VLAN and assigning the uplink port to it as an untagged member. Any ports that belong to the new VLAN are uplink network ports, carrying ACM traffic to and from the network.

Client VLANs

Client VLANs are special VLANs used by the module for client traffic. They have the following characteristics:

- Up to 24 client VLANs, depending on your configuration, may be used on a 5300xl switch. If two Access Controller xl Modules are installed in a 5300xl switch, the total number of VLANs used by the two modules may not exceed 24.
- Uplink network ports may not be members of a client VLAN.

When a port is added to a client VLAN the following changes are made to the port:

- Information used for ARP and MAC address processing is flushed.
- If GVRP is enabled, it is disabled and a message is displayed.
- If LACP passive is configured, it is disabled and a message is displayed.

Downlink client ports must be members of some other VLAN before they can be removed from a client VLAN. If you use the **no access-controller <slot-id> client-ports [e] <port-list>** command to remove an untagged downlink client port with no other VLAN memberships from a client VLAN, the port is automatically placed in the DEFAULT-VLAN as an untagged member. If you attempt to remove a tagged downlink client port that belongs to no other VLAN, the removal fails. Add the port to another VLAN, then delete it from the client VLAN.

Client VLANs can be configured without specifying any switch ports using the **access-controller <slot-id> client-ports vlan <vlan-list>** command from the configuration context. The VLANs are created with only the downlink port, **<slot-id>DP**, as a tagged member. Later you can use VLAN commands from the 5300xl CLI to add switch ports to this VLAN as downlink client ports.

Static VLAN Features Supported on Client VLANs

Client VLANs are special and they don't support all of the features of a regular 5300xl static VLAN. Table 12-9-1 below outlines the feature limitations of client VLANs.

Table 9-1. 5300xl Static VLAN Features on Client VLANs

5300xl Static VLAN Feature	Client VLAN Support
ACLs	Do not work. A warning issued.
IGMP	Not allowed.
IP Address	Not allowed.
IP Helper Address	Not allowed.
IRDP	Not allowed.
Management VLAN	A client VLAN may not be used.
Multicast routing	Not allowed.
OSPF	Not allowed.
PIM	Not allowed.
Primary VLAN	A client VLAN may not be used.
Protocol VLAN	A client VLAN may not be used.
RIP	Not allowed.
XRRP	Not allowed.

General Operating Rules

- Uplink and downlink ports cannot be members of the same VLAN.
- Switch 5300xl features used to manage ports that are connected to bridges don't apply, as the ACM is not a bridge.
- A client VLAN containing the downlink port, **<slot-id>DP**, is automatically created when the ACM is installed in a 5300xl switch. The VID for this VLAN is the **vlan-base** (default: 2000). You cannot remove a client VLAN if it is the only remaining VLAN with the downlink port as a member.
- Client VLANs may not be configured as the Management or Primary VLAN on the 5300xl switch.
- Multiple subnets on a downlink client port are not supported.
- Shut down the ACM
 - before resetting or reloading the 5300xl chassis
 - before turning off the 5300xl chassis.
 - before removing the module from the 5300xl chassis.

See the **shutdown** command in 12-“ACM Configuration Commands Summary and Syntax” on page 12-20.

Configuring the ACM on the Network

By default, the ACM uses DHCP to get an IP address. The uplink port of the ACM must be an untagged member of a VLAN that can communicate with the 740wl/760wl. If that communication is routed, the **Default gateway** needs to be present in the IP address configuration. When the IP address is assigned manually be sure to configure the **Default gateway** if it is needed. (See “Using the ACM's Extended CLI” on page 12-27.)

Use the following commands to configure an IP address manually.

Note

'ProCurve' is used as a generic prompt for all 5300xl switches. The term 'id' is used below for 'slot-id' to shorten the command prompt.

ProCurve (config)# access-controller <slot-id>

where **<slot-id>** is the slot in the 5300xl where the ACM is installed.

ProCurve (access-controller-id)# ip address <<ip-addr>/<1-32> <ip-addr> <mask>>

where **<ip-addr>/<1...32>** is the selected address in CIDR notation (/mask bit number), for example 10.1.2.3/24.

<ip-addr> <mask> provides the selected address and the mask.

If necessary, use the following command to set or change the default gateway:

ProCurve (access-controller-id)# ip default-gateway <ip-addr>

where **<ip-addr>** is the numeric IP address of the default gateway, for example 10.1.2.3.

Use the IP address of the 740wl/760wl and its shared secret to establish communications with the ACM.

ProCurve (access-controller-id)# access-control-server ip <ip-addr> secret <secret> <secret>

where **<ip-addr>** is the address of the 740wl/760wl

<secret> is the shared secret configured on the 740wl/760wl.

In the following example, an ACM establishes communications with an access control server with IP address 13.13.13.8. The access control server has a shared secret of 7734Oh. The **show status** command is used to confirm that communications has been established, indicated by a time value displayed (2 secs) in the Connected: field.

```
ProCurve Switch 5308xl (access-controller-B)#  
access-control-server ip 13.13.  
13.8 secret 77340h 77340h  
  
Shared secret changed.  
  
ProCurve Switch 5308xl (access-controller-B)# show status  
Uptime:          23 hrs, 12 mins  
Access Controller Function  
Access Control Server: 13.13.13.8  
Connected:       2 secs  
Active Clients:  None  
  
ProCurve Switch 5308xl (access-controller-B)#
```

Figure 12-4. Example of ACM Establishing Communication

Configuring the Access Controller xl Module

Once the module has an IP address and is communicating with its Access Control Server or Integrated Access Manager, configure downlink client ports, client VLANs, uplink network ports, and the uplink VLANs on the 5300xl switch.

Configuring Downlink Client Ports

Each downlink client port is automatically assigned to a unique client VLAN. The VID of the first client VLAN configured is specified by the **vlan-base** (default: 2000). Additional client VLANs use the next available sequential VID (2001, 2002, 2003, ...). If two Access Controller xl Modules are installed in the 5300xl switch, the **vlan-base** is the VID of the first client VLAN configured by either ACM. The next client VLAN configured on either ACM uses the next available sequential VID. Switch ports become untagged members of the client VLAN. The downlink port also becomes a tagged member of the client VLAN.

From the CLI command prompt at the global configuration level, enter

ProCurve (config) #access-controller <slot-id> client-ports <port list>

where

<slot-id> is the slot letter where the module is installed.

<port list> is the switch port(s) to be used as downlink client ports.

For example:

ProCurve (config)# access-controller b client-ports a2,a6

assigns two downlink client ports and two new client VLANs (see Figure 12-12-5). BDP, the downlink port for the module in slot B, is a tagged member of both client VLANs.

```
ProCurve Switch 5308xl(config)# access-controller b client-ports a2,a6
ProCurve Switch 5308xl(config)# access-controller b
ProCurve Switch 5308xl(access-controller-B)# show vlans
Downlink:
VLAN ID   VLAN Name   Ports
  2000   VLAN2000    A2,BDP
  2001   VLAN2001    A6,BDP

Uplink:
VLAN ID   VLAN Name   Ports
    1   DEFAULT_VLAN  A1,A3-A5,A7-A24,BUP

ProCurve Switch 5308xl(config)# show vlans 2000

Status and Counters - VLAN Information - Ports - VLAN 2000

802.1Q VLAN ID : 2000
Name : VLAN2000
Status : Port-based
Voice : No

Port Information Mode   Unknown VLAN Status
-----
A2           Untagged Disable   Down
BDP           Tagged  Disable     Up
```

Figure 12-5. Configuring A Downlink Client Port

**Notes on
Creating
Downlink Client
Ports**

Depending on how many VLANs are already configured in the 5300xl switch, the following messages may occur:

- Maximum Number of VLANs (X) has already been reached
Increase the maximum number of VLANs allowed on the switch.
- Command will take effect after saving configuration and reboot.
The switch requires a reboot to incorporate the new client VLANs into the system. If downlink client ports are added over a period of time, a reboot may be required after each addition to make the client VLAN available.
- Maximum number of client VLANs have been configured.
Operation failed.
The maximum number of client VLANs for this configuration has been reached. An existing client VLAN must be removed before the requested VLAN can be added.

Changing the VLAN-Base

When the ACM is installed in the 5300xl switch, a VLAN is created for the internal downlink port (<*slot-id*>DP). By default, this client VLAN is VLAN ID 2000, the **vlan-base**. You may change this using the following command.

ProCurve (Config)# access-controller vlan-base <2-4094>

where <2-4094> is the starting VLAN ID (VID) used when a client VLAN is configured.

The **vlan-base** is used by the <As-ls>[no] **access-controller <slot-id> client-ports [ethernet] < port-list >** command when it configures client VLANs for the specified switch ports. The very first port specified for use as a downlink client port becomes an untagged member of the **vlan-base** VID. Subsequent downlink client ports are assigned as untagged members of the next available sequential VID, beginning at the **vlan-base**.

Configuring Client VLANs

You may configure a client VLAN with a specific VID, containing just the downlink port as a tagged member. Later, you can add an untagged 5300xl port as a downlink client port to carry client traffic.

Use the following command to configure a client VLAN:

ProCurve (Config)# access-controller <slot-id> client-ports vlan <vlan-list>

where **<slot-id>** is the slot letter where the module is installed.
<vlan-list> is the VID for the desired client VLAN.

Configuring Uplink Network Ports

Uplink network ports are any switch ports that are members of the uplink VLAN, that is, the VLAN where the ACM's internal uplink port is an untagged member. By default, the uplink port (**<slot-id>UP**) is an untagged member of the DEFAULT-VLAN on the 5300xl switch. In this default configuration, all members of the DEFAULT-VLAN are uplink network ports.

Configuring the Uplink VLAN

To change the uplink VLAN, make the internal uplink port an *untagged* member of a new VLAN. Be sure that the new VLAN allows communication with the 740wl/760wl, or communications is lost.

ProCurve (Config)# vlan 25 untagged <slot-id>up

where **slot-id** is the 5300xl switch slot where the ACM module is installed.

This command configures a new uplink VLAN, VID 25, for the ACM module installed in slot **n**.

ACM Configuration Commands Summary and Syntax

Command	Page
Configuration Context	
access-controller <slot-id>	1220
[no] access-controller <slot-id> client-ports [e] <port-list>	1221
[no] access-controller <slot-id> client-ports vlan <vlan-list>	1222
access-controller <slot-id> reload	1222
access-controller <slot-id> shutdown	1222
access-controller vlan-base <2-4094>	1222
Access Controller Context	
access-control-server ip <ip addr> secret <secret> <secret>	1222
enable extended-commands	1223
exit	1223
[no] ip address <<ip-addr>/<1-32> <ip addr> <mask>>	1223
[no] ip default-gateway <ip-addr>	1223
[no] page	1223
terminal <length <2-1000> width <61-1920>>	1223

Configuration Context Command Syntax

Syntax: access-controller <slot-id>

*Changes the CLI to the access controller context for the access controller in **slot-id** (a - h). The **exit** command returns the CLI to the configuration context.*

Syntax: [no] access-controller <slot-id> client-ports [ethernet] < port-list >

*Assigns switch ports (**port-list**) to separate client VLANs for the access controller in **slot-id** (a - h). The ports are removed from all other VLANs. GVRP and LACP port provisioning are disabled.*

The client VLAN has the following port membership: the switch port, as an untagged member, and the ACM's downlink port (<slot-id>DP), as a tagged member.

*The **vlan-base** VID is configured when the ACM is installed in the switch. By default, this is VLAN 2000, whose only member is the downlink port. The very first time this command is used, the first switch port configured becomes a member of client VLAN **vlan-base** VID. For example, by default, the first time this command is used to assign a switch port to a client VLAN it becomes an untagged member of VLAN 2000. The next client VLAN configured takes the next available sequential VID, starting from the **vlan-base**.*

*Use the **no** form of the command to remove downlink client ports from ALL client VLANs associated with the module in <slot-id>. If the removed port has no other VLAN memberships it is automatically placed in the DEFAULT VLAN as an untagged member.*

*If a client VLAN that contained the removed port was configured using an **access-controller <slot-id> client-ports [ethernet] < port-list >** command and the ACM's downlink port is the only remaining member, the client VLAN is also removed.*

*The **no** form of this command does not remove a client VLAN configured using the **access-controller <slot-id> client-ports vlan < vlan-list >** command.*

Syntax: [no] access-controller <slot-id> client-ports vlan < vlan-list >

*Configures client VLANs with the **VIDs** given, containing only the downlink port, (<slot-id>DP), as a tagged member.*

*The **no** form can be used to remove client VLANs that were configured using the **access-controller <slot-id> client-ports vlan < vlan-list >** command and contain only the downlink port.*

Syntax: access-controller <slot-id> reload]

*Reboots the access controller in **slot-id** with the current software version.*

Syntax: access-controller <slot-id> shutdown]

*Halts the access controller in **slot-id**.*

Syntax: access-controller vlan-base <2-4094>]

*Sets the starting VLAN ID (VID) for client VLANs configured by the **access-controller <slot-id> client-ports < port-list >** or the **access-controller <slot-id> client-ports vlan < vlan-list >** commands. Valid VIDs are 2 - 4094.*

Access Controller Context Command Syntax

Syntax: access-control-server ip <ip addr> secret <secret> <secret>]

Specifies the numeric ip address and shared secret for the access control server (740wl or 760wl) that provides services to the ACM.

The secret must match the shared secret configured on the 740wl/760wl. If the shared secret contains spaces, enclose the secret in double quotes (" "). The secret must be entered twice, identically.

Syntax: enable extended-commands

Changes the CLI to the access controller extended commands context. A limited set of commands from the 720wl CLI is provided here. See “Using the ACM’s Extended CLI” for more information.

Syntax: exit

Leaves the access controller context and returns the CLI to the global configuration context.

Syntax: [no] ip address <<ip-addr>/<1-32> | <ip addr> <mask>>

*Statically configures the ip address and subnet mask for the ACM. The **no** form removes the fixed ip address and enables DHCP. (Default: DHCP)*

Syntax: [no] ip default-gateway <ip-addr>

Statically configures the numeric ip-addr of the default gateway. If DHCP is enabled, this command overrides the default gateway supplied by DHCP.

Syntax: [no] page

Turns on or off paging of the CLI command responses and CLI command stack replay (the up arrow function).

Syntax: terminal <length <2-1000> | width <61-1920>>

length <2-1000> *Sets the number of lines that are displayed between pauses when page is on. The initial setting is inherited from the previous context. This setting remains in effect when this context is exited.*

width <61-1920> *Sets the maximum line width to be used for CLI output or CLI input. Lines longer than the specified value wrap to the next line. The initial setting is inherited from the previous context. This setting remains in effect when this context is exited.*

Displaying Access Controller xl Status from the 5300xl CLI

Show commands are available in both the configuration context and the access controller context of the 5300xl CLI. These commands display ACM status and configuration.

ACM Display Commands Summary and Syntax

Command	Page
Configuration Context	
show access-controller <slot-id>	1225
[client-ports]	1225
[uplinks]	1225
show access-controller vlans	1225
show access-controller vlan-base	1225
Access Controller Context	
show ip	1226
show status	1226
show temperature	1226
show vlans	1225

Configuration Context Command Syntax

Syntax: show access-controller <slot-id>

Displays the following for the access controller in **slot-id** (a - h).

Versions	ACM version information for support staff.
Current status	BIOS error booted booting booting timed out configuration fault failed to boot halted initializing not responding performing self-test rebooting running self-test failed shutting down shut down (safe for removal)
Uplink MAC Address	ACM MAC address that appears in the 740wl/760wl Administrative Console as the System ID .

Syntax: show access-controller <slot-id> [client-ports]

Displays the downlink client ports for each client VLAN configured on the access controller in **slot-id** (a - h)

Syntax: show access-controller <slot-id> [uplinks]

Displays the uplink network ports configured on the access controller in **slot-id** (a - h)

Syntax: show access-controller vlans

Displays the **802.1Q VID** and **Name** of all configured client VLANs on the 5300xl switch. If two ACMs are installed, their client VLANs are presented in the list.

Syntax: show access-controller vlan-base

Displays the starting VLAN ID (VID) for client VLANs configured by the **access-controller <slot-id> client-ports <port-list>** or the **access-controller <slot-id> client-ports vlan <vlan-list>** commands.

Access Controller Context Command Syntax

Syntax: show ip

Displays the current IP configuration for the ACM, including:

Hostname:

Domain Name:

IP Address:

DHCP enabled:

Default gateway:

DHCP server:

DNS servers:

WINS servers:

Syntax: show status

Displays an overview of the ACM's status, including:

Uptime:

Access Control Server:

Connected:

Active Clients:

Syntax: show temperature

Displays the current temperature in degrees Celsius of the main processor of the ACM module.

Syntax: show vlans

Displays the **VLAN ID**, **VLAN Name**, and **Ports** for all VLANs associated with the ACM's uplink and downlink ports.

Managing the ACM

Once the module is installed and configured, most management tasks are done on the Access Control Server 740wl or Integrated Access Manager 760wl, using the Administrative Console. The Access Controller Module is managed in the same manner as a 720wl. For more information, see the *ProCurve Secure Access 700wl Series Management and Configuration Guide*, available on the CD shipped with the ACM, or from the ProCurve Networking Web site at www.procurve.com. (Click on **Technical support**, then **Product manuals (all)**).

Using the ACM's Extended CLI

A set of extended commands, available from the Access-Controller context, may be used to configure an ACM or display its status. These commands are a subset of the 720wl CLI commands. They are documented in the *ProCurve Secure Access 700wl Series Management and Configuration Guide*, available on the CD shipped with the ACM, or from the ProCurve Networking Web site at www.procurve.com. (Click on **Technical support**, then **Product manuals (all)**).

To access this context, type the following in the ACM context:

ProCurve(access-controller-id) # enable extended-commands

The command line prompt for the extended context is:

ProCurve(access-controller-id-ext)#

The available commands are listed below. Detailed descriptions are found in Appendix A, "Command Line Interface" in the *ProCurve Secure Access 700wl Series Management and Configuration Guide*.

Command
[no] ip address <<ip-addr>/<1-32> <ip-addr> <mask>>
[no] ip default-gateway <ip-addr>
[no] page
access-control-server ip <ip-addr> secret <secret> <secret>
add bridging <atalk wnmpp> <custom-bridging-string>
cancel upgrade

Command

Clear Commands

```
clear accesscontrolserver
clear dhcpserver
clear dns
clear domainname
clear gateway
clear hostname
clear logs
clear sharedsecret
clear upgradeproxy
delete bridging <atalk | wnmp> | <custom-bridging-string>
enable extended-commands
exit
factoryreset
get upgrade <url> <key> [mindowngrade | reboot | version]
help
logoff client <all | mac <mac-addr>>
reboot [downgrade | same | upgrade]
```

Set Commands

```
set accesscontrolserver <ip-addr>
set bridging <on | off>
set clientprobes <interval> <timeout>
set dhcp <on | off>
set dhcpserver <ip-addr>
set dns <primary-ip-addr> [<secondary-ip-addr>]
set domainname <domain>
set forwardipbroadcasts <all | none | on <port> | off <port> | <port>>
set gateway <ip-addr>
set hostname <host>
set ip <ip-addr> [<mask>] | <ip-addr>/<1-32>
set logopt addcat <all | error | info | none | session>
set logopt catlevel <error | info | none | session> <critical | major | minor | trivial | never>
set logopt cats <mask>
set logopt delcat <all | error | info | none | session>
```


Command

set logopt level <critical major minor trivial never> set logopt nofuncs <on off> set logopt noids <on off> set logopt oflags <mask> set logopt shorttrace <on off> set logopt string <logparam> set sharedsecret <secret> <secret> set upgradeproxy [<on off>] [host <ip-addr> [<port>]] [user <user> [<password>]]
--

Show Commands

show accesscontrolserver
show bridging
show client mac <mac-addr> [rights]
show clientprobes
show clients [mac <mac-addr>]
[sort <mac | ip | user | machine | port | sessions | idle>] [reverse]
show dhcpserver
show forwardipbroadcasts
show id
show ip
show logopt [level | cats | oflags | catlevel]
show logs [<critical | major | minor | trivial | never>] [cat <all | error | info>] [max <lines>]
[for <count> <seconds | minutes | hours | days | weeks | months>]
[search <quoted-text>] [reverse]
show natdhcp
show product
show serial
show sharedsecret
show status
show syslogserver
show temperature
show time
show upgrade
show upgradeproxy
show version
show vlans

Command
show vpn
terminal length <2..1000>
terminal width <61..1920>

Downloading New Software to the Module

New software is loaded through the Access Control Server or Integrated Access Manager using the Administrative Console.

Resetting the Module to Factory Defaults

The ACM may be returned to the factory default configuration using one of the following methods.

- Using the **factoryreset** command in the ACM's extended command context in the 5300xl CLI.
- Boot another, non-ACM module in the slot. This also removes the encoded ACM configuration information from the **show running** or **show config** command output.
- Use the **erase startup-config** command in the 5300xl CLI.
- Return the 5300xl chassis to its factory default configuration using the Reset and Clear keys on the front panel. (Refer to “Clear/Reset: Resetting to the Factory Default Configuration” in the Troubleshooting appendix of the *Management and Configuration Guide* for your switch.) This also resets the ACM.

Operating Notes

- Bridged protocols, such as Appletalk, are supported through a single downlink client port, whose client VLAN contains the downlink port as an *untagged* member. This must be configured manually on the switch. Each ACM may have one downlink client port configured to support bridged protocols.
- ProCurve recommends that a downlink client port be a member of only one client VLAN. Downlink client ports should not be members of any other VLANs, as this would allow access to unauthorized clients. If a downlink client port must be a member of another VLAN, configure filtering on the 740wl or 760wl to remove network traffic that might be sent to unauthorized ports.
- ProCurve recommends that you create a Management VLAN on your 5300xl switch. This secures the management of the 5300xl switch, allowing it to be managed only by members of the Management VLAN.
- Client-to-client communications is not possible through an ACM.
- ProCurve Manager does not support the ACM at this time. Support is expected later in 2005.

BIOS POST Event Log Messages

If a critical BIOS power on self test (POST) failure occurs when the ACM is inserted into a slot in a 5300xl chassis, the message below is posted to the Event Log. The 5300xl switch resets the ACM, up to two times (a total of three attempts to pass the POST). If the ACM fails three consecutive times, the module does not power on. The 5300xl switch can operate successfully if this occurs.

Slot <slot-id> Access Control Module Bios POST tests failed, Post bitmap = 0xFFFF

The POST error bitmap values are explained below.

<i>0x0001</i>	IDE failure.
<i>0x0002</i>	System memory failure.
<i>0x0004</i>	Shadow memory failure.
<i>0x0020</i>	Protected memory failure.
<i>0x0040</i>	CMOS not ready error.
<i>0x0100</i>	Periodic timer failure.
<i>0x0800</i>	Device configuration error.
<i>0x1000</i>	Memory configuration error.
<i>0x2000</i>	Non-volatile RAM failure.
<i>0x4000</i>	External or CPU cache failure.
<i>0x8000</i>	Level2 cache failure.

Port Trunking

Contents

Overview	13-2
Port Trunk Features and Operation	13-4
Trunk Configuration Methods	13-5
Menu: Viewing and Configuring a Static Trunk Group	13-9
CLI: Viewing and Configuring Port Trunk Groups	13-11
Using the CLI To View Port Trunks	13-11
Using the CLI To Configure a Static or Dynamic Trunk Group	13-14
Web: Viewing Existing Port Trunk Groups	13-17
Trunk Group Operation Using LACP	13-18
Default Port Operation	13-20
LACP Notes and Restrictions	13-21
Trunk Group Operation Using the “Trunk” Option	13-24
How the Switch Lists Trunk Data	13-25
Outbound Traffic Distribution Across Trunked Links	13-25

Overview

This chapter describes creating and modifying port trunk groups. This includes non-protocol trunks and LACP (802.3ad) trunks.

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port trunks	n/a	page 13-9	page 13-11	page 13-17
configuring a static trunk group	none	page 13-9	page 13-15	—
configuring a dynamic LACP trunk group	LACP passive	—	page 13-15	—

Port trunking allows you to assign up to eight physical links to one logical link (trunk) that functions as a single, higher-speed link providing dramatically increased bandwidth. This capability applies to connections between backbone devices as well as to connections in other network areas where traffic bottlenecks exist. A *trunk group* is a set of up to eight ports configured as members of the same port trunk. Note that the ports in a trunk group do not have to be consecutive. For example:

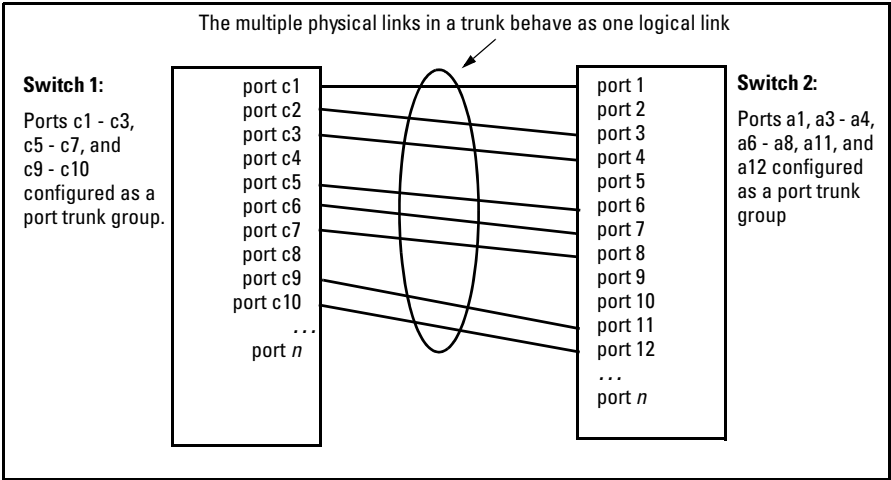


Figure 13-1. Conceptual Example of Port Trunking

With full-duplex operation in a eight-port trunk group, trunking enables the following bandwidth capabilities:

Port Connections and Configuration: All port trunk links must be point-to-point connections between a switch covered by this guide and another switch, router, server, or workstation configured for port trunking. No intervening, non-trunking devices are allowed. It is important to note that ports on both ends of a port trunk group must have the same mode (speed and duplex) and flow control settings.

Note

Link Connections. The switch does not support port trunking through an intermediate, non-trunking device such as a hub, or using more than one media type in a port trunk group. Similarly, for proper trunk operation, all links in the same trunk group must have the same speed, duplex, and flow control.

Port Security Restriction. Port security does not operate on a trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch resets the port security parameters for those ports to the factory-default configuration.

Caution

To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Port Trunk Features and Operation

The switches covered by this guide offer these options for port trunking:

- LACP: IEEE 802.3ad—page 13-18
- Trunk: Non-Protocol—page 13-24

The number of trunk groups supported on a given switch depends on the switch model and the number of ports physically available on the switch. The maximum theoretical bandwidths shown below are based on full-duplex operation.

Trunk Port Count	5300 1-Gig Links	3400/6400 Gig & 10-Gig	
		1-Gig Links	10-Gig Links
2	Up to 4 Gbs	Up to 4 Gbs	Up to 40 Gbs
3	6	6	60
4	8	8	80
5	8**	10	N/A*
6	8**	12	N/A*
7	8**	14	N/A*
8	8**	16	N/A*

*Although the 6400 can theoretically support an 8-port trunk, anything over 4 ports would not be practical because the trunk bandwidth capacity with 5 or more ports would exceed the bandwidth capacity of the remaining non-trunk ports. In any case, 80 Gbs is the theoretical maximum bandwidth for the 6400 switches.

** The maximum theoretical bandwidth for trunking on the 5300xl switches is 8 Gbs.

(Using the Link Aggregation Control Protocol—LACP—option, you can include standby trunked ports in addition to the maximum of eight actively trunking ports.)

LACP Note

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, etc.) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, ProCurve recommends that you leave the port Mode settings at **Auto** (the default). LACP also operates with **Auto-10**, **Auto-100**, and **Auto-1000** (if negotiation selects FDx), and **10FDx**, **100FDx**, and **1000FDx** settings. (The 10-gigabit ports available for some switch models allow only the **Auto** setting.)

Fault Tolerance: If a link in a port trunk fails, the switch redistributes traffic originally destined for that link to the remaining links in the trunk. The trunk remains operable as long as there is at least one link in operation. If a link is restored, that link is automatically included in the traffic distribution again. The LACP option also offers a standby link capability, which enables you to keep links in reserve for service if one or more of the original active links fails. See “Trunk Group Operation Using LACP” on page 13-18.)

Trunk Configuration Methods

Dynamic LACP Trunk: The switch automatically negotiates trunked links between LACP-configured ports on separate devices, and offers one dynamic trunk option: LACP. To configure the switch to initiate a dynamic LACP trunk with another device, use the **interface** command in the CLI to set the default LACP option to **Active** on the ports you want to use for the trunk. For example, the following command sets ports C1-C4 to LACP active:

```
ProCurve(config) int c1-c4 lacp active
```

Note that the preceding example works if the ports are not already operating in a trunk. To change the LACP option on ports already operating as a trunk, you must first remove them from the trunk. For example, if ports C1 - C4 were LACP-active and operating in a trunk with another device, you would do the following to change them to (the default) LACP-passive:

```
ProCurve(config)# no int c1-c4 lacp
```

Removes the ports from the trunk.

```
ProCurve(config)# int c1-c4 lacp passive
```

Configures LACP passive.

Static Trunk: The switch uses the links you configure with the Port/Trunk Settings screen in the menu interface or the **trunk** command in the CLI to create a static port trunk. The switch offers two types of static trunks: LACP and Trunk.

Table 13-1. Trunk Types Used in Static and Dynamic Trunk Groups

Trunking Method	LACP	Trunk
Dynamic	Yes	No
Static	Yes	Yes

Table 13-2. Trunk Configuration Protocols

Protocol	Trunking Options
LACP (802.3ad)	<p>Provides dynamic and static LACP trunking options.</p> <ul style="list-style-type: none">• Dynamic LACP — Use the switch-negotiated dynamic LACP trunk when:<ul style="list-style-type: none">– The port on the other end of the trunk link is configured for Active or Passive LACP.– You want fault-tolerance for high-availability applications. If you use an eight-link trunk you can also configure one or more additional links to operate as standby links that will activate only if another active link goes down.• Static LACP — Use the manually configured static LACP trunk when:<ul style="list-style-type: none">– The port on the other end of the trunk link is configured for a static LACP trunk– You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group.– <i>You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (Refer to “VLANs and Dynamic LACP” on page 13-23.)</i>– You want to use a monitor port on the switch to monitor an LACP trunk. <p>For more information, refer to “Trunk Group Operation Using LACP” on page 13-18.</p>
Trunk (non-protocol)	<p>Provides manually configured, static-only trunking to:</p> <ul style="list-style-type: none">• Most ProCurve switches and routing switches not running the 802.3ad LACP protocol.• Windows NT and HP-UX workstations and servers <p>Use the Trunk option when:</p> <ul style="list-style-type: none">– The device to which you want to create a trunk link is using a non-802.3ad trunking protocol– You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol.– You want to use a monitor port on the switch to monitor traffic on a trunk. <p>Refer to “Trunk Group Operation Using the “Trunk” Option” on page 13-24.</p>

Table 13-3. General Operating Rules for Port Trunks

Media: For proper trunk operation, all ports on both ends of a trunk group must have the same media type and mode (speed and duplex). (For the switches covered by this guide, ProCurve recommends leaving the port Mode setting at **Auto** or, in networks using Cat 3 cabling, **Auto-10**.)

Port Configuration: The default port configuration is **Auto**, which enables a port to sense speed and negotiate duplex with an Auto-Enabled port on another device. ProCurve recommends that you use the **Auto** setting for all ports you plan to use for trunking. Otherwise, you must manually ensure that the mode setting for each port in a trunk is compatible with the other ports in the trunk.

Recommended Port Mode Setting for LACP					
ProCurve(config)# show interface config					
Port Settings					
Port	Type	Enabled	Mode	Flow Ctrl	MDI
-----	-----	+	-----	-----	-----
C1	10/100TX	Yes	Auto	Disable	Auto
C2	10/100TX	Yes	Auto	Disable	Auto

Figure 13-2. Recommended Port Mode Setting for LACP

All of the following operate on a per-port basis, regardless of trunk membership:

- Enable/Disable
- Flow control (Flow Ctrl)

LACP is a full-duplex protocol. Refer to “Trunk Group Operation Using LACP” on page 13-18.

Trunk Configuration: All ports in the same trunk group must be the same trunk type (LACP or Trunk). All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.

A trunk appears as a single port labeled **Dyn1** (for an LACP dynamic trunk) or **Trk1** (for a static trunk of type: LACP, Trunk) on various menu and CLI screens. For a listing of which screens show which trunk types, see “How the Switch Lists Trunk Data” on page 13-25.

For spanning-tree or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for spanning-tree or VLAN operation.)

Traffic Distribution: All of the switch trunk protocols use the SA/DA (Source Address/Destination Address) method of distributing traffic across the trunked links. See “Outbound Traffic Distribution Across Trunked Links” on page 13-25.

Spanning Tree: 802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch). 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch). For each SpanningTree instance, you can adjust Spanning Tree parameters on a per-port basis. A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named **Trk1**, they are listed in the Spanning Tree display as **Trk1** and do not appear as individual ports in the Spanning Tree displays.

In this example showing part of the **show spanning-tree** listing, ports C1 and C2 are members of TRK1 and do not appear as individual ports in the port configuration part of the listing.

Port	Type	Cost	Priority	State	Designated Bridge
C3	100/1000T	5	128	Forwarding	0020c1-b27ac0
C4	100/1000T	5	128	Forwarding	0060b0-889e00
C5	100/1000T	5	128	Disabled	
C6	100/1000T	5	128	Disabled	
Trk1		1	64	Forwarding	0001e7-a0ec00

Figure 13-3. Example of a Port Trunk in a Spanning Tree Listing

When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.

Note: A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI **show spanning-tree** display, but not in the Spanning Tree Operation display of the Menu interface. If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.

IP Multicast Protocol (IGMP): A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN. A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or **show ip igmp** listing.

VLANs: Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.

Note: For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See “Trunk Group Operation Using LACP” on page 13-18.

Port Security: Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the **show port-security** listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you will see the following message and the command will not be executed:
< port-list > Command cannot operate over a logical port.

Monitor Port:

Note: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.

Menu: Viewing and Configuring a Static Trunk Group

Important

Configure port trunking *before* you connect the trunked links to another switch, routing switch, or server. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See “Using the CLI To Enable or Disable Ports and Configure Port Mode” on page 10-9.)

To View and/or Configure Static Port Trunking: This procedure uses the Port/Trunk Settings screen to configure a static port trunk group on the switch.

1. Follow the procedures in the Important note above.
2. From the Main Menu, Select:
2. Switch Configuration ...
2. Port/Trunk Settings
3. Press [E] (for **E**dit) and then use the arrow keys to access the port trunk parameters.

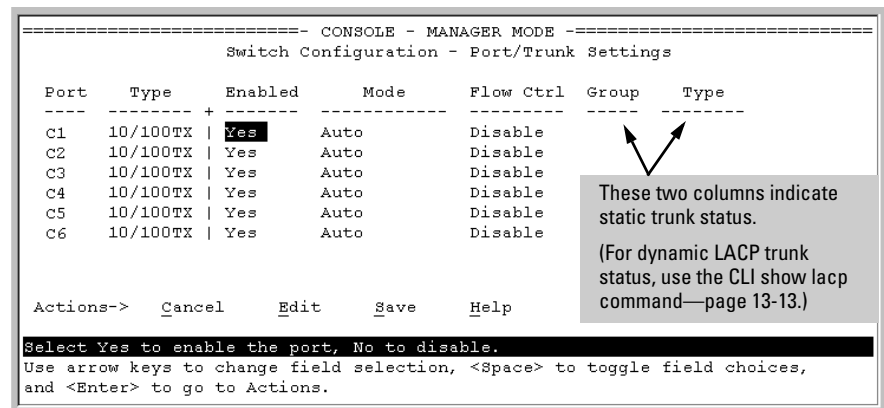


Figure 13-4. Example of the Menu Screen for Configuring a Port Trunk Group

4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group assignment (**Trk1**, **Trk2**, and so on) for the selected port.

Port Trunking

Menu: Viewing and Configuring a Static Trunk Group

- For proper trunk operation, all ports in a trunk must have the same media type and mode (such as 10/100TX set to 100FDx, or 100FX set to 100FDx). The flow control settings must also be the same for all ports in a given trunk. To verify these settings, see “Viewing Port Status and Configuring Port Parameters” on page 10-2.
- You can configure the trunk group with up to eight ports per trunk. If multiple VLANs are configured, all ports within a trunk will be assigned to the same VLAN or set of VLANs. (With the 802.1Q VLAN capability built into the switch, more than one VLAN can be assigned to a trunk. Refer to the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.)

(To return a port to a non-trunk status, keep pressing the Space bar until a blank appears in the highlighted Group value for that port.)

===== CONSOLE - MANAGER MODE =====						
Switch Configuration - Port/Trunk Settings						
Port	Type	Enabled	Mode	Flow Ctrl	Group	Type
C1	10/100TX	Yes	Auto	Disable		
C2	10/100TX	Yes	Auto	Disable		
C3	10/100TX	Yes	Auto	Disable		
C4	10/100TX	Yes	Auto	Disable		
C5	10/100TX	Yes	Auto	Disable	Trk1	Trunk
C6	10/100TX	Yes	Auto	Disable	Trk1	Trunk
Actions-> Cancel Edit Save Help						
Select whether the port is part of a trunk or Mesh.						
Use arrow keys to change field selection, <Space> to toggle field choices, and <Enter> to go to Actions.						

Figure 13-5. Example of the Configuration for a Two-Port Trunk Group

6. Move the cursor to the Type column for the selected port and use the Space bar to select the trunk type:
 - LACP
 - Trunk (the default type if you do not specify a type)

All ports in the same trunk group on the same switch must have the same Type (**LACP** or **Trunk**).

7. When you are finished assigning ports to the trunk group, press **[Enter]**, then **[S]** (for **Save**) and return to the Main Menu. (It is not necessary to reboot the switch.)

During the Save process, traffic on the ports configured for trunking will be delayed for several seconds. If the Spanning Tree Protocol is enabled, the delay may be up to 30 seconds.

8. Connect the trunked ports on the switch to the corresponding ports on the opposite device. If you previously disabled any of the trunked ports on the switch, enable them now. (See “Viewing Port Status and Configuring Port Parameters” on page 10-2.)

Check the Event Log (“Using the Event Log To Identify Problem Sources” on page C-27) to verify that the trunked ports are operating properly.

CLI: Viewing and Configuring Port Trunk Groups

Trunk Status and Configuration Commands

show trunks	below
show lacp	page 13-13
trunk	page 13-15
interface < port-list > lacp	page 13-15

Using the CLI To View Port Trunks

You can list the trunk type and group for all ports on the switch or for selected ports. You can also list LACP-only status information for LACP-configured ports.

Listing Static Trunk Type and Group for All Ports or for Selected Ports.

Syntax: show trunks [<port-list>]

Omitting the <port-list> parameter results in a static trunk data listing for all LAN ports in the switch. For example, in a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in figures 13-6 and 13-7 for displaying port data for ports belonging to static trunks.

Using a port list specifies, for switch ports in a static trunk group, only the ports you want to view. In this case, the command specifies ports A5 through A7. However, because port A6 is not in a static trunk group, it does not appear in the resulting listing:

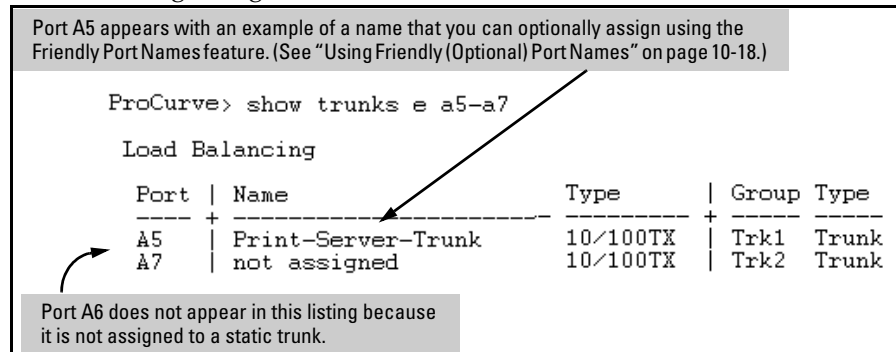


Figure 13-6. Example Listing Specific Ports Belonging to Static Trunks

The **show trunks <port-list>** command in the above example includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In figure 13-7, the command does not include a port list, so the switch lists all ports having static trunk membership.

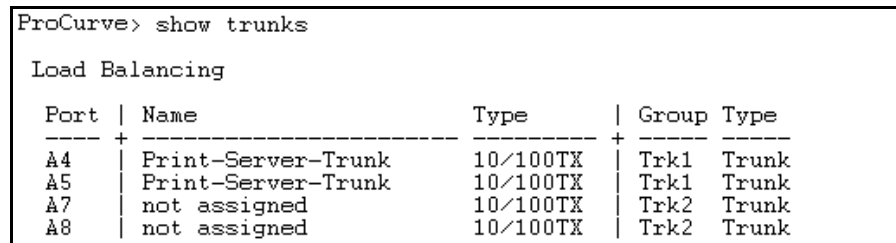


Figure 13-7. Example of a Show Trunk Listing Without Specifying Ports

Listing Static LACP and Dynamic LACP Trunk Data.

Syntax: show lacp

Lists data for only the LACP-configured ports..

In the following example, ports A1 and A2 have been previously configured for a static LACP trunk. (For more on “Active”, see table 11-13-5 on page 13-21.)

ProCurve> show lacp					
LACP					
PORT	LACP	TRUNK	PORT	LACP	LACP
NUMB	ENABLED	GROUP	STATUS	PARTNER	STATUS
----	-----	-----	-----	-----	-----
A1	Active	Trk1	Up	Yes	Success
A2	Active	Trk1	Up	Yes	Success
A3	Active	A3	Down	No	Success
A4	Passive	A4	Down	No	Success
A5	Passive	A5	Down	No	Success
A6	Passive	A6	Down	No	Success

Figure 13-8. Example of a Show LACP Listing

(For a description of each of the above-listed data types, refer to table 13-5, “LACP Port Status Data” on page 13-21.)

Dynamic LACP Standby Links. Dynamic LACP trunking enables you to configure standby links for a trunk by including more than eight ports in a dynamic LACP trunk configuration. When eight ports (trunk links) are up, the remaining link(s) will be held in standby status. If a trunked link that is “Up” fails, it will be replaced by a standby link, which maintains your intended bandwidth for the trunk. (See also the “Standby” entry under “Port Status” in "Table 13-5. LACP Port Status Data" on page 13-21.) In the next example, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows Standby status, while the remaining eight links are “Up”.

ProCurve> show lacp						
LACP						
	PORT NUMB	LACP ENABLED	TRUNK GROUP	PORT STATUS	LACP PARTNER	LACP STATUS
"Up" Links	A1	Active	Dyn1	Up	Yes	Success
	A2	Active	Dyn1	Up	Yes	Success
	A3	Active	Dyn1	Up	Yes	Success
	A4	Active	Dyn1	Up	Yes	Success
	A5	Active	Dyn1	Up	Yes	Success
	A6	Active	Dyn1	Up	Yes	Success
	A7	Active	Dyn1	Up	Yes	Success
	A8	Active	Dyn1	Up	Yes	Success
Standby Link	A9	Active	Dyn1	Standby	Yes	Success

Figure 13-9. Example of a Dynamic LACP Trunk with One Standby Link

Using the CLI To Configure a Static or Dynamic Trunk Group

Important

Configure port trunking *before* you connect the trunked links between switches. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See “Using the CLI To Enable or Disable Ports and Configure Port Mode” on page 10-9.)

The table on page 13-4 describes the maximum number of trunk groups you can configure on the switches covered in this guide. An individual trunk can have up to eight links, with additional standby links if you’re using LACP. You can configure trunk group types as follows:

Trunk Type	Trunk Group Membership	
	TrkX (Static)	DynX (Dynamic)
LACP	Yes	Yes
Trunk	Yes	No

Note

Trunks configured as FEC (Fast Ethernet Channel) are not supported. To configure port trunk groups, use static or LACP trunks. For release notes describing the latest software updates, visit the ProCurve Networking web site at www.procurveve.com. Click on **Technical support**, and then click on **Product manuals (all)**.

The following examples show how to create different types of trunk groups.

Configuring a Static Trunk or Static LACP Trunk Group.

Syntax: `trunk <port-list> <trk1 ... trk36> <trunk | lacp>`

Configures the specified static trunk type.

This example uses ports C4 - C6 to create a non-protocol static trunk group with the group name of **Trk2**.

```
ProCurve(config)# trunk c4-c6 trk2 trunk
```

Removing Ports from a Static Trunk Group. This command removes one or more ports from an existing **Trkx** trunk group.

Caution

Removing a port from a trunk can create a loop and cause a broadcast storm. When you remove a port from a trunk where spanning tree is not in use, ProCurve recommends that you first disable the port or disconnect the link on that port.

Syntax: `no trunk <port-list>`

Removes the specified ports from an existing trunk group.

For example, to remove ports C4 and C5 from an existing trunk group.

```
ProCurve(config)# no trunk c4-c5
```

Enabling a Dynamic LACP Trunk Group. In the default port configuration, all ports on the switch are set to LACP **Passive**. However, to enable the switch to automatically form a trunk group that is dynamic on both ends of the link, the ports on one end of a set of links must be LACP **Active**. The ports on the other end can be either LACP **Active** or LACP **Passive**. The **active** command enables the switch to automatically establish a (dynamic) LACP trunk group when the device on the other end of the link is configured for LACP **Passive**.

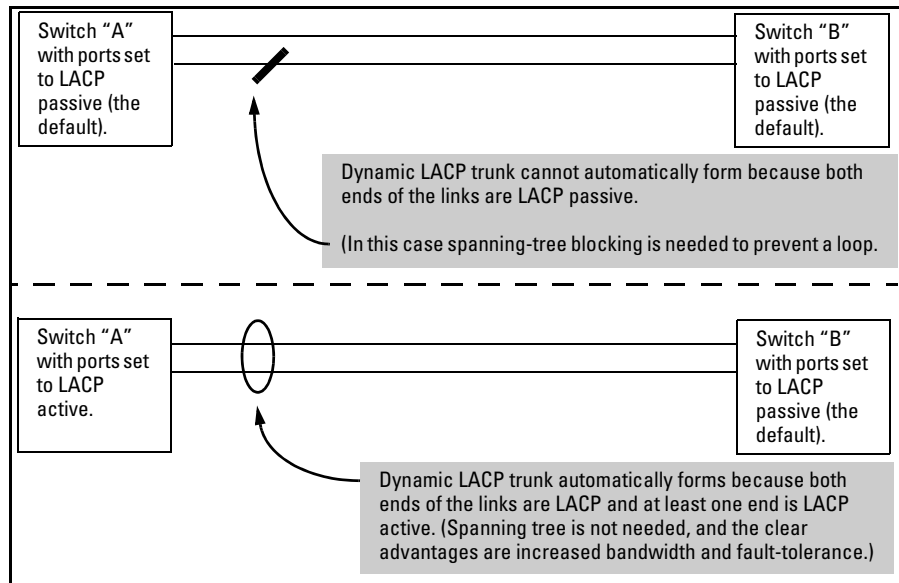


Figure 13-10. Example of Criteria for Automatically Forming a Dynamic LACP Trunk

Syntax: `interface <port-list> lacp active`

Configures <port-list> as LACP active. If the ports at the other end of the links on <port-list> are configured as LACP passive, then this command enables a dynamic LACP trunk group on <port-list>.

This example uses ports C4 and C5 to enable a dynamic LACP trunk group.

```
ProCurve(config)# interface c4-c5 lacp active
```

Removing Ports from an Dynamic LACP Trunk Group. To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP **Active** and LACP **passive** without first removing LACP operation from the port.)

Caution

Unless spanning tree is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where spanning tree is not in use, ProCurve recommends that you first disable the port or disconnect the link on that port.

Syntax:

Syntax: no interface <port-list> lacp

Removes <port-list> from any dynamic LACP trunk and returns the ports in <port-list> to passive LACP.

In this example, port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, you would do the following:

```
ProCurve(config)# no interface c6 lacp
ProCurve(config)# interface c6 lacp passive
```

Note that in the above example, if the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Web: Viewing Existing Port Trunk Groups

While the web browser interface does not enable you to configure a port trunk group, it does provide a view of an existing trunk group.

To view any port trunk groups:

Click on the **Status** tab.

Click on **[Port Status]**.

Trunk Group Operation Using LACP

The switch can automatically configure a dynamic LACP trunk group or you can manually configure a static LACP trunk group.

Note

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, etc.) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, ProCurve recommends that you leave the port Mode settings at **Auto** (the default). LACP also operates with **Auto-10**, **Auto-100**, and **Auto-1000** (if negotiation selects FDx), and **10FDx**, **100FDx**, and **1000FDx** settings.

LACP trunk status commands include:

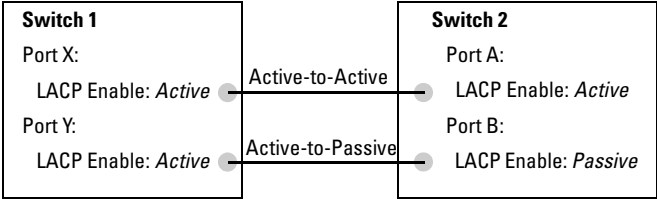
Trunk Display Method	Static LACP Trunk	Dynamic LACP Trunk
CLI show lacp command	Included in listing.	Included in listing.
CLI show trunk command	Included in listing.	Not included.
Port/Trunk Settings screen in menu interface	Included in listing.	Not included

Thus, to display a listing of dynamic LACP trunk ports, you must use the **show lacp** command.

In most cases, trunks configured for LACP on the switches covered by this manual operate as described in table 13-4 on the next page.

Table 13-4. LACP Trunk Types

LACP Port Trunk Configuration	Operation
Dynamic LACP	<p>This option automatically establishes an 802.3ad-compliant trunk group, with LACP for the port Type parameter and DynX for the port Group name, where X is an automatically assigned value from 1 to 36, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 36 trunk groups in any combination of static and dynamic trunks.)</p> <p>Note: Dynamic LACP trunks operate only in the default VLAN (unless GVRP is enabled and Forbid is used to prevent the trunked ports from joining the default VLAN). Thus, if an LACP dynamic port forms using ports that are not in the default VLAN, the trunk will automatically move to the default VLAN unless GVRP operation is configured to prevent this from occurring. In some cases, this can create a traffic loop in your network. For more on this topic, refer to “VLANs and Dynamic LACP” on page 13-23.</p> <p>Under the following conditions, the switch automatically establishes a dynamic LACP port trunk group and assigns a port Group name:</p> <ul style="list-style-type: none">• The ports on both ends of each link have compatible mode settings (speed and duplex).• The port on one end of each link must be configured for LACP Active and the port on the other end of the same link must be configured for either LACP Passive (the default) or LACP Active. For example:



Either of the above link configurations allow a dynamic LACP trunk link.

The following operations are supported:

- Dynamic LACP <Active | Passive > to Static LACP
- Dynamic LACP Active to Dynamic LACP Active

NOT supported is:

- Dynamic LACP <Active | Passive> to <Trunk>

Backup Links: A maximum of eight operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more additional (backup) links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing eight-port dynamic LACP trunk, ensure that the ports in the standby link are configured as either active-to-active or active-to-passive between switches.

Displaying Dynamic LACP Trunk Data: To list the configuration and status for a dynamic LACP trunk, use the CLI **show lacp** command.

Note: The dynamic trunk is automatically created by the switch, and is not listed in the static trunk listings available in the menu interface or in the CLI **show trunk** listing.

LACP Port Trunk Configuration	Operation
Static LACP	<p>Provides a manually configured, static LACP trunk to accommodate these conditions:</p> <ul style="list-style-type: none">• A static LACP trunk will work with a dynamic LACP trunk. The VLAN membership of a dynamic trunk will be VLAN 1; the static LACP trunk should also be a member of VLAN 1. (Static trunks can be configured to be a member of another VLAN.)• You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group.• You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (Refer to “VLANs and Dynamic LACP” on page 13-23.)• You want to use a monitor port on the switch to monitor an LACP trunk. <p>The following operations are supported:</p> <ul style="list-style-type: none">• Dynamic LACP <Active Passive > to Static LACP• Dynamic LACP Active to Dynamic LACP Active <p>NOT supported is:</p> <ul style="list-style-type: none">• Dynamic LACP <Active Passive> to <Trunk> <p>(The table on page 13-4 lists the maximum number of trunk groups allowed on switches covered by this guide.)</p> <p>Displaying Static LACP Trunk Data: To list the configuration and status for a static LACP trunk, use the CLI show lacp command. To list a static LACP trunk with its assigned ports, use the CLI show trunk command or display the menu interface Port/Trunk Settings screen.</p> <p>Static LACP does not allow standby ports.</p>

Default Port Operation

For the Series 4200vl switches, LACP is turned off by default. For the Series 5300xl, 3400lcl, and 6400cl switches, all ports are configured for passive LACP by default.

When LACP is turned on, if LACP is not configured as Active on at least one end of a link, then the port does not try to detect a trunk configuration and operates as a standard, untrunked port. Table 13-5 lists the elements of per-port LACP operation. To display this data for a switch, execute the following command in the CLI:

```
ProCurve> show lacp
```


Table 13-5. LACP Port Status Data

Status Name	Meaning
Port Numb	Shows the physical port number for each port configured for LACP operation (C1, C2, C3 .). Unlisted port numbers indicate that the missing ports are assigned to a static Trunk group are not configured for any trunking.
LACP Enabled	<p>Active: The port automatically sends LACP protocol packets.</p> <p>Passive: The port does not automatically send LACP protocol packets, and responds only if it receives LACP protocol packets from the opposite device.</p> <p>A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports will not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device.</p> <p>Note: For the Series 4200vl switches, LACP is turned off by default. For the 5300xl, 3400cl, and 6400cl switches, all ports are configured for passive LACP by default.</p>
Trunk Group	<p>TrkX: This port has been manually configured into a static LACP trunk.</p> <p>Trunk Group Same as Port Number: The port is configured for LACP, but is not a member of a port trunk.</p>
Port Status	<p>Up: The port has an active LACP link and is not blocked or in Standby mode.</p> <p>Down: The port is enabled, but an LACP link is not established. This can indicate, for example, a port that is not connected to the network or a speed mismatch between a pair of linked ports.</p> <p>Disabled: The port cannot carry traffic.</p> <p>Blocked: LACP, spanning tree has blocked the port. (The port is not in LACP Standby mode.) This may be due to a (brief) trunk negotiation or a configuration error such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support. (See the table on page 13-4.)</p> <p>Standby: The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the Dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or “standby” unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a Standby port, if available, to replace the failed port.</p>
LACP Partner	<p>Yes: LACP is enabled on both ends of the link.</p> <p>No: LACP is enabled on the switch, but either LACP is not enabled or the link has not been detected on the opposite device.</p>
LACP Status	<p>Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link.</p> <p>Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore not able to send LACP packets across the link. This can be caused, for example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.</p>

LACP Notes and Restrictions

802.1x (Port-Based Access Control) Configured on a Port. To maintain security, LACP is not allowed on ports configured for 802.1x authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables 802.1x on that port.

```
ProCurve(config)# aaa port-access authenticator b1
LACP has been disabled on 802.1x port(s).
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port access (802.1x) is enabled. For example:

```
ProCurve(config)# int b1 lacp passive
Error configuring port < port-number >: LACP and 802.1x
cannot be run together.
ProCurve(config)#
```

To restore LACP to the port, you must first remove the port's 802.1x configuration and then re-enable LACP active or passive on the port.

Port Security Configured on a Port. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
ProCurve(config)# port-security a17 learn-mode static
address-limit 2
LACP has been disabled on secured port(s).
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
ProCurve(config)# int a17 lacp passive
Error configuring port A17: LACP and port security cannot
be run together.
ProCurve(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Changing Trunking Methods. To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP Trunks. Where a port is configured for LACP (Active or Passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

Dynamic LACP Trunks. You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the **trunk** command. (Refer to “Using the CLI To Configure a Static or Dynamic Trunk Group” on page 13-14.)

VLANs and Dynamic LACP. A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use **Forbid** to prevent the ports from joining the default VLAN).

- If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.
- If there are ports that you do not want on the default VLAN, ensure that they cannot become dynamic LACP trunk members. Otherwise a traffic loop can unexpectedly occur. For example:

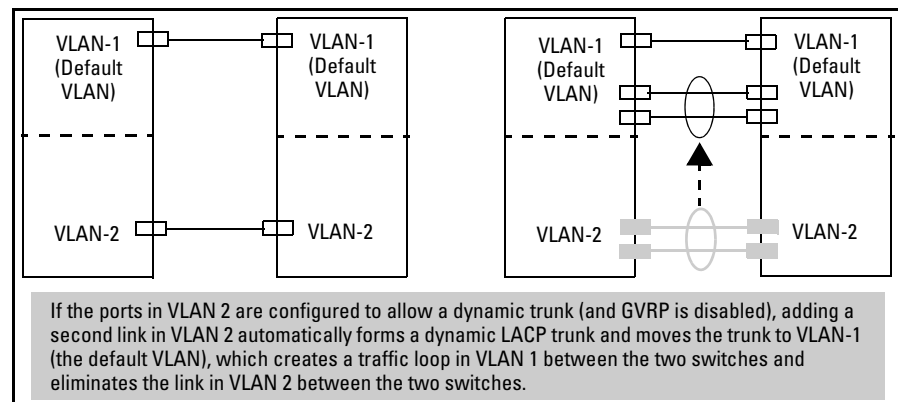


Figure 13-11. A Dynamic LACP Trunk Forming in a VLAN Can Cause a Traffic Loop

Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP. If Spanning Tree and/or IGMP is enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-Duplex and/or Different Port Speeds Not Allowed in LACP Trunks. The ports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx). The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking. (10-gigabit ports operate only at FDx.)

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx). However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If the port is a 10-gigabit port.
- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/Static LACP Interoperation: A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links will be ignored.

Trunk Group Operation Using the “Trunk” Option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

When a trunk group is configured with the **trunk** option, the switch automatically sets the trunk to a priority of “4” for spanning-tree operation (even if spanning-tree is currently disabled. This appears in the running-config file as `spanning-tree Trkn priority 4`. Executing **write memory** after configuring the trunk places the same entry in the startup-config file.

Use the Trunk option to establish a trunk group between a 5300xl, 3400cl, or 6400cl switch and another device, where the other device’s trunking operation fails to operate properly with LACP trunking configured on the 5300xl or LACP trunking configured on the 3400/6400cl switches or 4200vl switches.

How the Switch Lists Trunk Data

Static Trunk Group: Appears in the menu interface and the output from the CLI **show trunk** and **show interfaces** commands.

Dynamic LACP Trunk Group: Appears in the output from the CLI **show lacp** command.

Interface Option	Dynamic LACP Trunk Group	Static LACP Trunk Group	Static Non-Protocol (5300xl Switches Only)
Menu Interface	No	Yes	Yes
CLI show trunk	No	Yes	Yes
CLI show interfaces	No	Yes	Yes
CLI show lacp	Yes	Yes	No
CLI show spanning-tree	No	Yes	Yes
CLI show igmp	No	Yes	Yes
CLI show config	No	Yes	Yes

Outbound Traffic Distribution Across Trunked Links

The two trunk group options (LACP and Trunk) use source-destination address pairs (SA/DA) for distributing outbound traffic over trunked links.

SA/DA (source address/destination address) causes the switch to distribute outbound traffic to the links within the trunk group on the basis of source/destination address pairs. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and sends traffic from the same source address to a different destination address through a different link, depending on the rotation of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through different links. Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while others in the same trunk have unused bandwidth capacity even though the address assignments are evenly distributed across the links in a trunk. In actual networking environments, this is rarely a problem. However, if it becomes a problem, you can use the ProCurve

Manager Plus network management software to quickly and easily identify the sources of heavy traffic (top talkers) and make adjustments to improve performance.

Broadcasts, multicasts, and floods from different source addresses are distributed evenly across the links. As links are added or deleted, the switch redistributes traffic across the trunk group. For example, in figure 13-12 showing a three-port trunk, traffic could be assigned as shown in table 13-6.

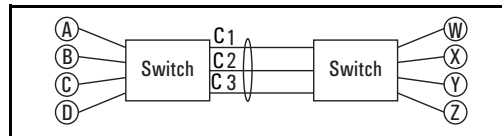


Figure 13-12. Example of Port-Trunked Network

Table 13-6. Example of Link Assignments in a Trunk Group (SA/DA Distribution)

Source:	Destination:	Link:
Node A	Node W	1
Node B	Node X	2
Node C	Node Y	3
Node D	Node Z	1
Node A	Node Y	2
Node B	Node W	3

Port Traffic Controls

Contents

Overview	14-3
All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches	14-4
Introduction	14-4
Rate-Limiting Operation	14-5
Configuring Inbound Rate-Limiting	14-5
Displaying the Current Rate-Limit Configuration	14-6
Operating Notes for Rate-Limiting	14-7
ICMP Rate-Limiting	14-10
Terminology	14-10
Effect of ICMP Rate-Limiting	14-10
Operating Notes for ICMP Rate-Limiting	14-17
Guaranteed Minimum Bandwidth (GMB) on the Series 5300xl Switches	14-21
Introduction	14-21
Terminology	14-21
GMB Operation	14-21
Configuring Guaranteed Minimum Bandwidth for Outbound Traffic	14-23
Displaying the Current Guaranteed Minimum Bandwidth Configuration	14-25
GMB Operating Notes	14-26
Jumbo Packets on the Series 3400cl and Series 6400cl Switches	14-27
Terminology	14-27
Operating Rules	14-28
Configuring Jumbo Packet Operation	14-29
Overview	14-29

Viewing the Current Jumbo Configuration	14-29
Enabling or Disabling Jumbo Traffic on a VLAN	14-31
Operating Notes for Jumbo Traffic-Handling	14-32
Troubleshooting	14-34

Overview

Feature	Default	Menu	CLI	Web
Rate-Limiting	None	n/a	14-4	n/a
Guaranteed Minimum Bandwidth	Per Queue: 8%-16%-30%-45%	n/a	14-21	n/a
Jumbo Packets (3400cl and 6400cl Only)	Disabled	n/a	14-27	n/a

This chapter includes:

- **Rate Limiting:** Enables a port to limit the amount of bandwidth a user or device may utilize for inbound traffic on the switch.
- **Guaranteed Minimum Bandwidth (GMB):** Provides a method for ensuring that each of a port’s outbound queues has a specified minimum consideration for sending traffic out on the link to another device.
- **Jumbo Packets (3400cl and 6400cl Only):** Enables ports operating at 1 Gbs or 10 Gbps speeds to accept inbound packets of up to 9220 bytes when configured for jumbo traffic.

All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches

Feature	Default	Menu	CLI	Web
rate-limit < limit-% >	none	n/a	page 14-5	n/a
show rate-limit [port-list]	n/a	n/a	page 14-6	n/a

Note

This feature applies to the 5300xl, 3400cl, and 6400cl switches.

Introduction

Rate-Limiting for all traffic provides a method for limiting the amount of bandwidth a user or device may utilize inbound on a switch port. This effectively sets an inbound usage level on a given port, and is a tool for enforcing maximum service level commitments granted to network users. This feature operates on a per-port level and is not configurable on port trunks. Note that rate-limiting is designed to be applied at the network edge to limit inbound traffic from non-critical users or to enforce service agreements such as those offered by Internet Service Providers (ISPs) to provide only the bandwidth for which a customer has paid.

Note

On the 5300xl switches beginning with software release E.09.02, rate-limiting also can be applied by a RADIUS server during an authentication client session. See the chapter “*RADIUS Authentication and Accounting*” in the *Security Guide*, version January 2005 or later.

Note

The 5300xl switches also support IGMP rate-limiting. For further information, see the chapter “*Multimedia Traffic Control with IP Multicast (IGMP)*” in the *Advanced Traffic Management Guide* for your switch.

Caution

Rate-Limiting is intended for use on edge ports in a network. It is not recommended for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

Under network stress conditions, a port may allow occasional bursts of inbound traffic forwarding that exceed the port's configured rate. For this reason, rate-limiting should not be employed as a security feature.

Rate-Limiting Operation

Rate-Limiting operates on a per-port basis to allow only the specified percentage of the port's bandwidth to be used for inbound traffic. For example, if a 100 Mbps port negotiates a link at 100 Mbps and is rate-limit configured at 50%, then the inbound traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows no more than 5 Mbps of inbound traffic.

Note for Rate-Limiting on Series 3400cl and 6400cl Switches

Configuring rate-limiting on a 3400cl or 6400cl switch port consumes one per-port rule and one per-port QoS mask. This affects the resources available for configuring QoS and ACLs. If you plan to configure QoS and/or ACLs on a 3400cl or 6400cl switch, refer to the chapters on these topics in the *Advanced Traffic Management Guide* for your switch.

Configuring Inbound Rate-Limiting

This command controls inbound usage of a port by setting a limit on the bandwidth available for inbound traffic. Beginning with software release E.10.02, the syntax of the rate-limiting command has changed to accommodate the new ICMP rate-limiting feature available. Refer to "ICMP Rate-Limiting" on page 14-10.

Syntax: [no] int < port-list > rate-limit < all | icmp > < 0..100 >

*Configures an inbound traffic rate limit (on non-trunked ports) as a percentage of the bandwidth available on the link. You can configure a rate limit from either the global configuration level (as shown above) or from the port context level. The "no" form of the command disables rate-limiting on the specified ports. (Default: **Disabled.**)*

Notes:

- Rate-Limiting applies only to non-trunked ports (and is not recommended for meshed ports).
- Configuring a rate limit of 0 (zero) on a port blocks all inbound traffic on that port. However, if this is the desired behavior for the port, ProCurve recommends that you use < **port-list** > **disable** to disable the port instead of configuring a rate limit of 0.

For example, either of the following commands configures an inbound rate limit of 60% on ports A3 - A5:

```
ProCurve (config)# int a3-a5 rate-limit all 60
ProCurve (eth-A3-A5)# rate-limit all 60
```

Displaying the Current Rate-Limit Configuration

This command displays the per-port rate-limit configuration in the running-config file.

Syntax: show rate-limit all [port-list]

Without [port-list], this command lists the rate-limit configuration for all ports on the switch. With [port-list], this command lists the rate-limit configuration for the specified port(s). This command operates the same way in any CLI context.

For example, if you wanted to view the rate-limiting configuration on the first five ports in the module in slot “A”:

```
ProCurve# show rate-limit all a1-a5
```

Inbound Rate Limit Maximum %	
Port	Total
A1	Disabled
A2	Disabled
A3	60
A4	60
A5	60

Figure 14-1. Example of Listing the Rate-Limit Configuration

The **show config** command lists the per-port rate-limiting and Guaranteed Minimum Bandwidth configuration in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

```
ProCurve(config)# show config status
Running configuration is same as the startup configuration.
ProCurve(config)# show config

Startup configuration:

; J4850A Configuration Editor; Created on release #E.08.00

hostname "HPswitch"
module 1 type J4820A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  bandwidth-min output 10 20 20 50
  exit
interface A2
  bandwidth-min output 10 20 20 50
  exit
interface A3
  rate-limit all 60
  exit
interface A4
  rate-limit all 60
  exit
interface A5
  rate-limit all 60
  exit
```

The outbound port priority queues 1 - 4 for ports A1-A2 are configured with the indicated Guaranteed Minimum Bandwidth percentages.

Ports A3-A5 are configured with a rate limit of 60 %. (Ports A1 and A2 are not configured for rate-limiting.)

Figure 14-2. Example of Rate-Limit Settings Listed in the “show config” Output

Operating Notes for Rate-Limiting

- Rate-Limiting is available on all types of ports on the switches covered by this guide, and at all port speeds configurable for these devices. (Rate-Limiting is not allowed on trunked ports.)
- The configured rate limit on a port reflects the permitted forwarding rate from the port to the switching fabric, and is visible as the *average* rate of the outbound traffic originating from the rate-limited port. Also, rate-limiting reflects the available percentage of a port's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from a rate-limited port to a particular queue of an outbound port are not measures of the actual rate limit enforced on a port.
- Rate-Limiting operates on a per-port basis, regardless of traffic priority. Configuring rate-limiting on a port where other features affect inbound port queue behavior (such as flow control) can result in the port not achieving its configured rate-limiting maximum. For example, in some situations with flow control configured on a rate-limited port, there can be enough “back pressure” to hold high-priority inbound traffic from the upstream device or application to a rate that is lower than the configured

rate limit. In this case, the inbound traffic flow does not reach the configured rate and lower priority traffic is not forwarded into the switch fabric from the rate-limited port. (This behavior is termed “head-of-line blocking” and is a well-known problem with flow-control.) In another type of situation, an outbound port can become oversubscribed by traffic received from multiple rate-limited ports. In this case, the actual rate for traffic on the rate-limited ports may be lower than configured because the total traffic load requested to the outbound port exceeds the port’s bandwidth, and thus some requested traffic may be held off on inbound.

**Note on Testing
Rate-Limiting**

Rate-Limiting is byte-based and is applied to the available bandwidth on a port, and not to any specific applications running through the port. If the total bandwidth requested by all applications together is less than the available, configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing software applications, as well as most regular network applications. Consider the following example, which uses the minimum packet size:

The total available bandwidth on a 100 Mbps port “X” (allowing for Inter-packet Gap—IPG), with no rate-limiting restrictions, is:

$$(((100,000,000 \text{ bits}) / 8) / 84) \times 64 = 9,523,809 \text{ bytes per second}$$

where:

- *The divisor (84) includes the 12-byte IPG, 8-byte preamble, and 64-bytes of data required to transfer a 64-byte packet on a 100 Mbps link.*
- *Calculated “bytes-per-second” includes packet headers and data. This value is the maximum “bytes-per-second” that 100 Mbps can support for minimum-sized packets.*

Suppose port “X” is configured with a rate limit of 50% (4,761,904 Mbytes). If a throughput-testing application is the only application using the port, and transmits 1 Mbyte of data through the port, it uses only 10.5% of the port’s available bandwidth, *and the rate-limit of 50% has no effect*. This is because the maximum rate permitted (50%) exceeds the test application’s bandwidth usage (126,642-164,062 bytes, depending upon packet size, which is only 1.3-1.7% of the available total). Before rate-limiting can occur, the test application’s bandwidth usage must exceed the configured rate-limit. In this example, the bandwidth usage must exceed 50% of the port’s total available bandwidth. That is, in this example, to test the rate-limit setting, the following must be true:

$$\text{bandwidth usage} > (0.50 \times 9,523,809)$$

- **Network Stress Conditions:** Under normal network operating conditions, rate-limiting limits inbound traffic on a port to no more than the configured level. However, under network stress conditions, the port may allow occasional bursts of inbound traffic forwarding that exceed the configured rate.
- **Optimum Rate-Limiting Operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Outbound Traffic Flow:** Configuring rate-limiting on a port does not control the rate of outbound traffic flow on the port.
- **Rate-Limiting Effect on Port Trunks:** Rate-Limiting is not supported on ports configured in a trunk group. Configuring a port for rate-limiting and then adding it to a trunk suspends rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message:

`<port-list>: Operation is not allowed for a trunked port.`
- **Traffic Filters on Rate-Limited Ports:** Configuring a traffic filter on a port does not prevent the switch from including filtered traffic in the bandwidth-use measurement for rate-limiting. That is, where rate-limiting and traffic filtering are configured on the same port, the inbound, filtered traffic is included in the bandwidth measurement for calculating when the limit has been reached. Traffic filters include:
 - ACLs
 - Source-Port filters
 - Protocol filters
 - Multicast filters
- **Rate-Limiting Not Recommended on Mesh Ports:** Rate-Limiting can reduce the efficiency of paths through a mesh domain.

ICMP Rate-Limiting

In IP networks, ICMP messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers). Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network. ICMP Rate-Limiting provides a method for limiting the amount of bandwidth that may be utilized for inbound ICMP traffic on a switch port or trunk. This feature allows users to restrict ICMP traffic to levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect). In addition, this preserves inbound port bandwidth for non-ICMP traffic.

Terminology

All-Traffic Rate-Limiting: Applies a rate-limit to all inbound traffic, including ICMP traffic, received on an interface.

ICMP Rate-Limiting: Applies a rate-limit to all inbound ICMP traffic received on an interface, but does not limit other types of inbound traffic.

Spoofed Ping: An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.

Effect of ICMP Rate-Limiting

ICMP rate-limiting generally allows only a specified percentage of an interface's inbound bandwidth to be used for ICMP traffic. As a result, inbound bandwidth is preserved for non-ICMP traffic and the port or trunk throttles any sudden flood of inbound ICMP traffic that may be due to a worm or virus attack (or any other cause). Notice that ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can configure both ICMP rate-limiting and all-traffic rate-limiting.

Caution

The ICMP protocol is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior, and should normally be configured to allow one to five per cent of available inbound bandwidth to be used for ICMP traffic. ***This feature should not be used to remove all ICMP traffic from a network.***

Note

Because all-traffic rate-limiting and ICMP rate-limiting operate similarly, the CLI command for the all-traffic version of rate-limiting has been modified for compatibility with the ICMP rate-limiting CLI command. Beginning with software release E.10.02, these commands appear in the following format:

rate-limit [all | icmp]

For more on all-traffic rate-limiting, refer to the chapter titled “Port Traffic Controls” in the *Management and Configuration Guide* for your switch (January 2005 or greater).

Network Application. Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. (“Normal” ICMP traffic levels should be the maximums that occur when the network is rebooting.)

Port Traffic Controls

All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches

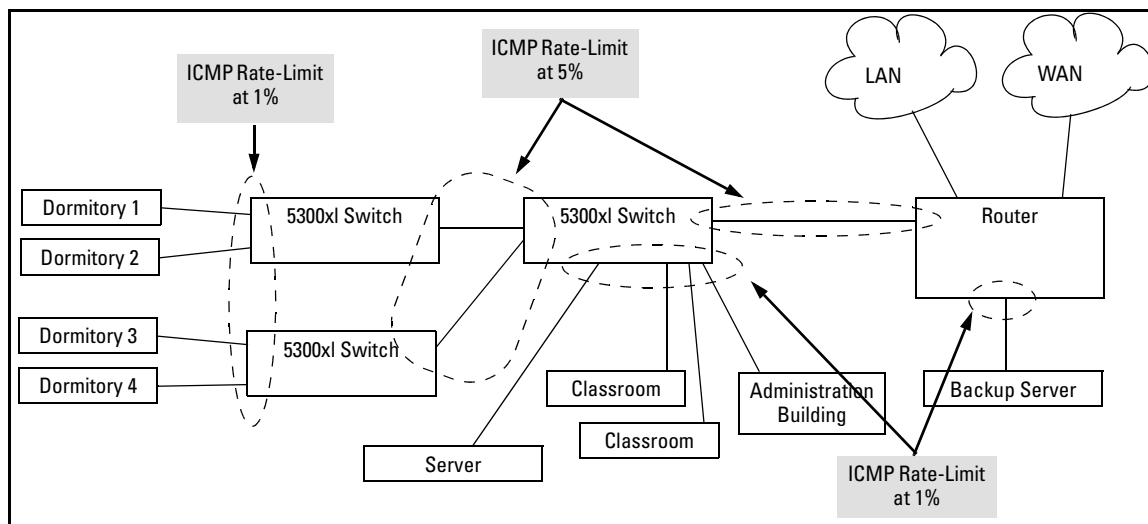


Figure 14-3. Example of ICMP Rate-Limiting

ICMP Rate-Limiting Operation. ICMP rate-limiting operates on an interface (per-port or per-trunk) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic. For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured at 5%, then the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows 0.5 Mbps of inbound traffic. (For more on performance under varying operating conditions, refer to “Operating Notes for ICMP Rate-Limiting” on page 14-17.) If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured).

Applying ICMP Rate-Limiting to a Port Trunk. These rules apply to ICMP rate-limiting when applied to a trunk:

- The configured ICMP traffic limit is applied as a percentage of all traffic inbound on the trunk.
- ICMP rate-limiting is only supported on a port trunk where all members of the trunk are *in the same module slot*. ICMP rate-limiting is **not** supported on trunks having members in multiple module slots.

- All ports belonging to a trunk configured for ICMP rate-limiting operate according to the trunk configuration, regardless of the ICMP rate-limiting state that existed on the port prior to its being added to the trunk. (While a port is in a trunk, any ICMP rate-limiting previously configured for that port is suspended, but remains in the switch configuration.)
- Removing a port from a trunk returns the port to whatever ICMP rate-limiting state existed on the port before it was put into the trunk.

Note

A rate-limited trunk should include only ports on the same slot/module. A rate-limited trunk configured across module boundaries is not supported and produces unpredictable rate-limiting operation and results.

Using Both ICMP and All-Traffic Rate-Limiting on an Interface. ICMP and all-traffic rate-limiting can be configured on the same interface. All-Traffic rate-limiting applies to all inbound traffic (including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.

Note that if the inbound, all-traffic load on an interface meets or exceeds the current all-traffic rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, then all excess traffic will be dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has been reached). Suppose, for example:

- The all-traffic limit on port “X” is configured at 55% of the port’s bandwidth.
- The ICMP traffic limit on port “X” is configured at 2% of the port’s bandwidth.

If at a given moment:

- inbound ICMP traffic on port “X” is using 1% of the port’s bandwidth, and
- inbound traffic of all types on port “X” demands 61% of the port’s bandwidth,

then all inbound traffic above 55% of the port’s bandwidth, including any additional ICMP traffic will be dropped as long as all inbound traffic combined on the port demands 55% or more of the port’s bandwidth.

Note

Under network stress conditions, an interface may allow occasional bursts of inbound ICMP traffic forwarding that exceed the interface’s configured ICMP traffic rate. Refer to “ICMP Rate-Limit Imposes an Average Bandwidth Limit” on page 18.

Configuring Inbound Rate-Limiting. This command controls inbound usage of a port by setting a limit on the bandwidth available for inbound traffic.

Syntax: [no] int < port-list | trunk-list > rate-limit icmp < 0..100 >

*Configures inbound ICMP traffic rate limiting. You can configure a rate limit from either the global configuration level (as shown above) or from the interface context level. The **no** form of the command disables ICMP rate-limiting on the specified interface(s). (Default: **Disabled**.)*

1 - 99: *Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.*

0 : *This value causes an interface to drop all incoming ICMP traffic, and is not recommended. Refer to the Caution on page 11.*

Note: *ICMP Rate-Limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain).*

For example, either of the following commands configures an inbound rate limit of 1% on ports A3 - A5, which are used as network edge ports:

```
ProCurve(config)# int a3-a5 rate-limit icmp 1
ProCurve (eth-A3-A5)# rate-limit icmp 1
```

Displaying the Current Rate-Limit Configuration. This command displays the per-interface rate-limit configuration in the running-config file.

Syntax: show rate-limit icmp [port-list | trunk-list]

*Without [**port-list** | **trunk-list**], this command lists the ICMP rate-limit configuration for all ports or trunks on the switch. With [**port-list** | **trunk-list**], this command lists the rate-limit configuration for the specified interface(s). This command operates the same way in any CLI context.*

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot “B”:

```
ProCurve(config)# show rate-limit icmp b1-b6
```

Inbound ICMP Rate Limit Maximum Percentage		
Port	Rate	Limit
B1	Disabled	
B2	1	
B3	1	
B4	1	
B5	1	
B6	Disabled	

Ports B2-B5 are configured with an ICMP rate limit of 1%. (Ports B1 and B6 are not configured for ICMP rate-limiting.)

Figure 14-4. Example of Listing the Rate-Limit Configuration

The **show running** command displays the currently applied setting for any interfaces in the switch configured for ICMP rate limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

Port Traffic Controls

All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches

```
ProCurve(config)# show config status
Running configuration is same as the startup configuration.

ProCurve(config)# show config
Startup configuration:

; J4850A Configuration Editor; Created on release #E.10.00

hostname "HPswitch"
module 2 type J8161A
module 4 type J8161A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged B1-B24,D1-D24
  ip address dhcp-bootp
  exit
access-controller vlan-base 2000
interface B2
  rate-limit icmp 1
  exit
interface B3
  rate-limit icmp 1
  exit
interface B4
  rate-limit icmp 1
  exit
interface B5
  rate-limit icmp 1
  exit
```

The **show config status** command compares the content of the startup-config and running-config files and prints a report.

Ports B2-B5 are configured with an ICMP rate limit of 1%.

Figure 14-5. Example of ICMP Rate-Limit Settings Listed in the “show running” Output

ICMP Rate-Limiting Trap and Event Log Messages. If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of the when the event occurred on the port.) For example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded
configured limit on port 1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the following **setmib** command.

Syntax: `setmib hplcmpRatelimitPortAlarmflag.< internal-port-#> -i 1`

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

For example, an operator noticing an ICMP rate-limiting trap or Event Log message originating with port A1 on a 5300xl switch would use the following **setmib** command to reset the port to send a new message if the condition occurs again.

```
ProCurve(config)# setmib hpicmpratelimitportalarm-  
flag.1 -i 1
```

Operating Notes for ICMP Rate-Limiting

Note on Testing Rate-Limiting

ICMP rate-limiting is byte-based and is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, then no ICMP rate-limit can be applied. That is, an interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both **rate-limit all** and **rate-limit icmp**, then the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit it is necessary to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, it is necessary to ensure that the ICMP traffic volume exceeds the configured maximum. Note also that testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP average aggregate traffic over time.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

- **Interface Support:** ICMP rate-limiting is available on all types of ports and trunks on the switches covered by this guide, and at all port speeds configurable for these devices.
- **Rate-Limiting Not Permitted on Mesh Ports:** Either type of rate-limiting can reduce the efficiency of paths through a mesh domain.
- **Monitoring (Mirroring) ICMP Rate-Limited Interfaces:** If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by “drop” or “forward” decisions.)
- **ICMP Rate-Limit Imposes an Average Bandwidth Limit:** The configured ICMP rate limit on an interface reflects the permitted *average* forwarding rate for ICMP traffic from the interface to the switching fabric. (Note that while occasional bursts of traffic above the configured rate may be observed, the average rate will conform to the configured limit). Rate-Limiting is packet-based, and is calculated internally as the maximum number of 64-byte packets that can be forwarded within the configured bandwidth percentage. Where traffic includes packets larger than 64 bytes, actual average rates may be lower than the configured rate. Also, ICMP rate-limiting reflects the available percentage of an interface’s entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.
- **Network Stress Conditions:** Under normal network operating conditions, ICMP rate-limiting limits inbound traffic on an interface to no more than the configured level. However, under network stress conditions, the interface may allow occasional, brief bursts of inbound traffic forwarding that exceed the configured rate.
- **Below-Maximum Rates:** ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough “back pressure” to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed “head-of-line blocking” and is a well-known problem with flow-control.) In cases where both types of rate-limiting (**rate-limit all** and **rate-limit icmp**) are configured on the same interface, this situation is more likely to occur. In another type of situa-

tion, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited interfaces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.

- **Optimum Rate-Limiting Operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Heavy Memory Usage:** Combinations of intensive QoS, rate-limiting, and/or IDM ACL service demands on a switch can impose heavy memory usage on the switch's dynamic hardware rule processor, and can sometimes result in slower system performance. In such cases, moving support for some of the service load to other devices can improve performance.
- **Outbound Traffic Flow:** Configuring ICMP rate-limiting on an interface does not control the rate of outbound traffic flow on the interface.
- **Traffic Filters on Rate-Limited Interfaces:** Configuring a traffic filter on an interface does not prevent the switch from including filtered traffic in the bandwidth-use measurement for either type of rate-limiting (ICMP or all). That is, where rate-limiting and traffic filtering are configured on the same interface, the inbound, filtered traffic is included in the bandwidth measurement for calculating when the limit has been reached. Traffic filters include:
 - ACLs
 - Source-Port filters
 - Protocol filters
 - Multicast filters
- **Determining the 5300xl Switch Port Number Used in ICMP Port Reset Commands To Enable Excess ICMP Traffic Notification Traps and Event Log Messages:** Use the internal port numbers described in this section with the **setmib** command described on page 17. The port number included in the command corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number, use the **walkmib ifDescr** command, as shown in the following figure:

Port Traffic Controls

All-Traffic Rate-Limiting for the 5300xl, 3400cl and 6400cl Switches

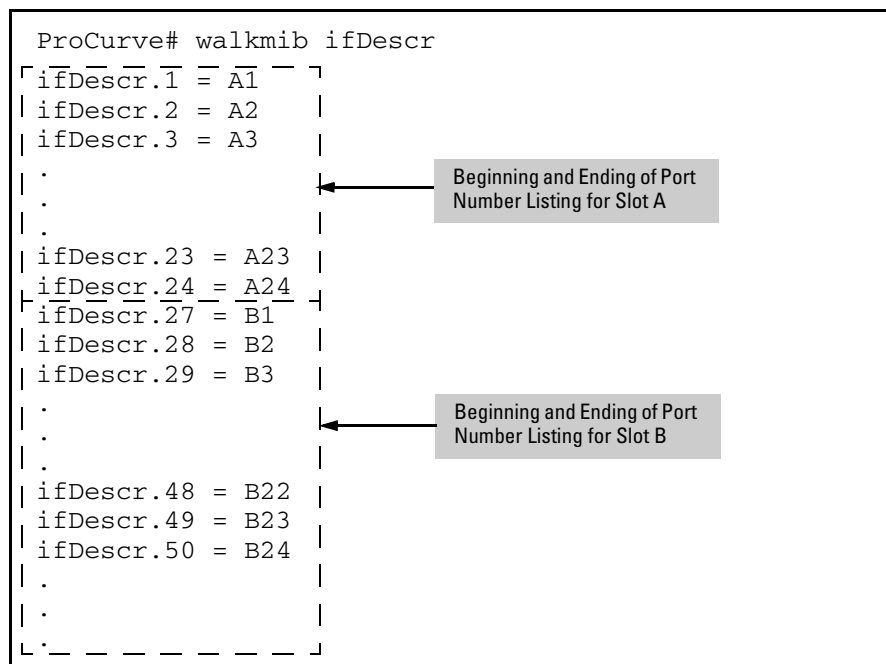


Figure 14-6. Matching Internal Port Numbers to External Slot/Port Numbers

Guaranteed Minimum Bandwidth (GMB) on the Series 5300xl Switches

This section applies only to the Series 5300xl switches.

Feature	Default	Menu	CLI	Web
bandwidth-min output	Per-Queue: 8%-16%-30%-45%	n/a	page 14-23	n/a
show bandwidth output [port-list]	n/a	n/a	page 14-21	n/a

Introduction

Guaranteed Minimum Bandwidth (GMB) provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum consideration for sending traffic out on the link to another device. This can prevent a condition where applications generating lower-priority traffic in the network are frequently or continually "starved" by high volumes of higher-priority traffic. You can configure GMB per-port or per-trunk.

Terminology

Oversubscribed Queue: The condition where there is insufficient bandwidth allocated to a particular outbound priority queue for a given port. If additional, unused bandwidth is not available, the port delays or drops the excess traffic.

GMB Operation

The switch services per-port outbound traffic in a descending order of priority; that is, from the highest priority to the lowest priority. Each port offers four prioritized, outbound traffic queues. Tagged VLAN traffic is prioritized according to the 802.1p priority the traffic carries. Untagged VLAN traffic is assigned a priority of "0" (normal).

Table 14-1. Per-Port Outbound Priority Queues

802.1p Priority Settings in Tagged VLAN Packets*	Outbound Priority Queue for a Given Port
1 (low)	1
2 (low)	
0 (normal)	2
3 (normal)	
4 (medium)	3
5 (medium)	
6 (high)	4
7 (high)	

*The switch processes outbound traffic from an untagged port at the "0" (normal) priority level.

You can use GMB to reserve a specific percentage of each port's available outbound bandwidth for each of the four priority queues. This means that regardless of the amount of high priority outbound traffic on a port, you can ensure that there will always be bandwidth reserved for lower-priority traffic.

Since the switch services outbound traffic according to priority (highest to lowest), the highest-priority outbound traffic on a given port automatically receives the first priority in servicing. Thus, in most applications, it is necessary only to specify the minimum bandwidth you want to allocate to the lower three priority queues. In this case, the high-priority traffic automatically receives all unassigned bandwidth without starving the lower-priority queues.

Conversely, configuring a bandwidth minimum on only the high-priority outbound queue of a port (and not providing a bandwidth minimum for the lower-priority queues) is not recommended because it may "starve" the lower-priority queues. (See the Note, below.)

Note

For a given port, when the demand on one or more outbound queues exceeds the minimum bandwidth configured for those queues, the switch apportions unallocated bandwidth to these queues on a priority basis. As a result, specifying a minimum bandwidth for a high-priority queue but not specifying a minimum for lower-priority queues can starve the lower-priority queues during periods of high demand on the high priority queue. For example, if a port configured to allocate a minimum bandwidth of 80% for outbound high-priority traffic experiences a demand above this minimum, then this burst starves lower-priority queues that *do not have a minimum configured*. Normally, this will not altogether halt lower priority traffic on the network, but will likely cause delays in the delivery of the lower-priority traffic.

The sum of the GMB settings for all four outbound queues on a given port cannot exceed 100%.

Configuring Guaranteed Minimum Bandwidth for Outbound Traffic

For any port or group of ports you can configure either the default minimum bandwidth settings for each outbound priority queue or a customized bandwidth allocation. For most applications, ProCurve recommends configuring GMB with the same values on all ports on the switch so that the outbound traffic profile is consistent for all outbound traffic. However, there may be instances where it may be advantageous to configure special profiles on connections to servers or to the network infrastructure (such as links to routers, other switches, or to the network core).

Syntax: [no] int < port-list > bandwidth-min output

Configures the minimum bandwidth allocation for the outbound priority queue for each port in < port-list >. The default values are:

- Queue 1 (low priority): 8%
- Queue 2 (normal or unmarked priority): 16%
- Queue 3 (medium priority): 30%
- Queue 4 (high priority): 45%

*The **no** form of the command disables GMB for all ports in < port-list >. In this state, which is the equivalent of setting all outbound queues on a port to 0 (zero), a high level of higher-priority traffic can starve lower-priority queues, which can slow or halt lower-priority traffic in the network. You can configure bandwidth minimums from either the global configuration level (as shown above) or from the port context level. For information on outbound port queues, refer to table 14-1, “Per-Port Outbound Priority Queues” on page 14-22.*

Syntax: [no] int < port-list > bandwidth-min output (*Continued*)

[< queue1% > < queue2% > < queue3% > < queue4% >]

*For ports in < port-list >, specifies the minimum outbound bandwidth as a percent of the total bandwidth for each outbound queue. The queues receive service in descending order of priority. You must specify a bandwidth percent value for all four queues, and the sum of the bandwidth percentages must not exceed 100%. (0 is a value for a queue percentage setting. See the **Note**, below.) Configuring a total of less than 100% across the four queues results in unallocated bandwidth that remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority. For example, if you configure a minimum of 10% for queues 1 - 3, and 0% for queue 4, then the unallocated bandwidth will be available to all four queues in the following prioritized order:*

- 1. Queue 4 (high priority)*
- 2. Queue 3 (medium priority)*
- 3. Queue 2 (normal priority)*
- 4. Queue 1 (low priority)*

A setting of 0 (zero %) on a queue means that no bandwidth minimum is specifically reserved for that queue for each of the ports in < port-list >. Also, there is no benefit to setting the high-priority queue (queue 4) to 0 (zero) unless you want the medium queue (queue 3) to be able to support traffic bursts above its guaranteed minimum.

Notes: *Configuring 0% for a queue can result in that queue being starved if any higher queue becomes oversubscribed and is then given all unused bandwidth.*

The switch applies the bandwidth calculation to the link speed the port is currently using. For example, if a 10/100 Mbs port negotiates to 10 Mbps on the link, then it bases its GMB calculations on 10 Mbps; not 100 Mbps.

*Use **show bandwidth output < port-list >** to display the current GMB configuration. (The **show config** and **show running** commands do not include GMB configuration data.)*

For example, suppose you wanted to configure the following outbound minimum bandwidth availability for ports A1 and A2:

Priority of Outbound Port Queue	Minimum Bandwidth %	Effect on Outbound Bandwidth Allocation
4	50	Queue 4 has the first priority use of all outbound bandwidth not specifically allocated to queues 1 - 3. If, for example, bandwidth allocated to queue 1 is not being used and queues 3 and 4 become oversubscribed, queue 4 has first-priority use of the unused bandwidth allocated to queue 1.
3	20	Queue 3 has a guaranteed minimum bandwidth of 20% available for outbound traffic. If queue 3 becomes oversubscribed and queue 4 is not already using all of the unallocated bandwidth, then queue 3 can use the unallocated bandwidth. Also, any unused bandwidth allocated to queues 1 or 2 is available to queue 3 if queue 4 has not already claimed it.
2	20	Queue 2 has a guaranteed minimum bandwidth of 20% and, if oversubscribed, is subordinate to queues 4 and 3 in priority for any unused outbound bandwidth available on the port.
1	10	Queue 1 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues 4, 3, and 2 for any unused outbound bandwidth available on the port.

Either of the following commands configures ports A1 and A2 with the bandwidth settings shown in the preceding table:

```
ProCurve(config)#int a1-a2 bandwidth-min output 10 20 20 50
ProCurve(eth-A1-A2)#bandwidth-min output 10 20 20 50
```

Displaying the Current Guaranteed Minimum Bandwidth Configuration

This command displays the per-port GMB configuration in the running-config file.

Syntax: show bandwidth output [port-list]

Without [port-list], this command lists the GMB configuration for all ports on the switch. With [port-list], this command lists the GMB configuration for the specified ports. This command operates the same way in any CLI context. If the command lists Disabled for a port, there are no bandwidth minimums configured for any queue on the port. (Refer to the description of the no form of the bandwidth-min output command on page 14-23.)

For example, to display the GMB configuration resulting from either of the above commands:

```
ProCurve# show bandwidth output a1-a5
```

Outbound Guaranteed Minimum Bandwidth %

Port	Low Priority	Normal Priority	Medium Priority	High Priority
A1	10	20	20	50
A2	10	20	20	50
A3	8	16	30	45
A4	8	16	30	45
A5	8	16	30	45

Annotations:

- Arrows point from the 'User-Configured Minimum Bandwidth Settings' label to the values for ports A1 and A2.
- An arrow points from the 'Default Minimum Bandwidth Settings' label to the values for ports A3, A4, and A5.

Figure 14-7. Example of Listing the Guaranteed Minimum Bandwidth Configuration

For an example listing the GMB configuration in the startup-config file, refer to figure 14-2 on page 14-7.

GMB Operating Notes

Granularity of Applied GMB Settings. Incremental bandwidth settings greater than 0 and less than 100 are internally computed in steps of 1.6%. Thus, the switch internally converts a configured bandwidth percentage to the closest multiple of 1.6.

Jumbo Packets on the Series 3400cl and Series 6400cl Switches

This section applies only to the ProCurve Series 3400cl and Series 6400cl switches.

Feature	Default	Menu	CLI	Web
display VLAN jumbo status	n/a	—	14-29	—
configure jumbo VLANs	Disabled	—	14-31	—

The *Maximum Transmission Unit* (MTU) is the maximum size IP packet the switch can receive for Layer 2 packets inbound on a port. The switch drops any inbound packets larger than the MTU allowed on the port. On ports operating at 10 Mbps or 100 Mbps, the MTU is fixed at 1522 bytes. However, ports operating at 1 Gbs or 10 Gbps speeds accept forward packets of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. In the 3400cl/6400cl switches you can enable inbound jumbo packets on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and *operating* at 1 Gbs or 10 Gbps allow inbound jumbo packets of up to 9220 bytes. (Regardless of the mode configured on a given jumbo-enabled port, if the port is operating at only 10 Mbps or 100 Mbps, only packets that do not exceed 1522 bytes are allowed inbound on that port.)

Terminology

Jumbo Packet: On the 3400cl/6400cl switches, an IP packet exceeding 1522 bytes in size. The maximum Jumbo packet size is 9220 bytes. (This size includes 4 bytes for the VLAN tag.)

Jumbo VLAN: A VLAN configured to allow inbound jumbo traffic. All ports belonging to a jumbo and operating at 1 Gbps or higher can receive jumbo packets from external devices. If the switch is in a meshed domain, then all meshed ports (operating at 1 Gbps or higher) on the switch will accept jumbo traffic from other devices in the mesh.

MTU (*Maximum Transmission Unit*): This is the maximum-size IP packet the switch can receive for Layer 2 packets inbound on a port. The switch allows jumbo packets of up to 9220 bytes.

Standard MTU: On the 3400cl/6400cl switches, an IP packet of 1522 bytes in size. (This size includes 4 bytes for the VLAN tag.)

Operating Rules

- **Required Port Speed:** The 3400cl/6400cl switches allow inbound and outbound jumbo packets on ports operating at speeds of 1 gigabit or higher. At lower port speeds, only standard (1522-byte or smaller) packets are allowed, regardless of the jumbo configuration.
- **Flow Control:** Disable flow control (the default setting) on any ports or trunks through which you want to transmit or receive jumbo packets. Leaving flow control enabled on a port can cause a high rate of jumbo drops to occur on the port.
- **Switch Meshing:** If you enable jumbo traffic on a VLAN in a 3400cl or 6400cl switch, then all meshed ports on the switch will be enabled to support jumbo traffic. (On a given meshed switch, every meshed port operating at 1 Gbps or higher becomes a member of every VLAN configured on the switch.)
- **GVRP Operation:** A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.
- **Port Adds and Moves:** If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.
- **Jumbo Traffic Sources:** A port belonging to a jumbo-enabled VLAN can receive inbound jumbo packets through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, then port 1 can receive jumbo traffic from devices on either VLAN. For a method to allow only some ports in a VLAN to receive jumbo traffic, refer to “Operating Notes for Jumbo Traffic-Handling” on page 14-32.

Configuring Jumbo Packet Operation

Command	Page
show vlans	14-30
show vlans ports < port-list >	14-31
show vlans < vid >	14-31
jumbo	14-31

Overview

1. Determine the VLAN membership of the ports or trunks through which you want the switch to accept inbound jumbo traffic. For operation with GVRP enabled, refer to the GVRP topic under “Operating Rules”, above.
2. Ensure that the ports through which you want the switch to receive jumbo packets are operating at least at gigabit speed. (Check the **Mode** field in the output for the **show interfaces brief < port-list >** command.)
3. Use the **jumbo** command to enable jumbo packets on one or more VLANs statically configured in the switch. (All ports belonging to a jumbo-enabled VLAN can receive jumbo packets.
4. Execute **write memory** to save your configuration changes to the startup-config file.

Viewing the Current Jumbo Configuration

Syntax: show vlans

*Lists the static VLANs configured on the switch and includes a **Jumbo** column to indicate which VLANs are configured to support inbound jumbo traffic. All ports belonging to a jumbo-enabled VLAN can receive jumbo traffic. (For more information refer to “Operating Notes for Jumbo Traffic-Handling” on page 14-32.) See figure 14-8, below.*

ProCurve(config)# show vlans

Status and Counters - VLAN Information

Maximum VLANs to support : 8

Primary VLAN : DEFAULT_VLAN

Management VLAN :

802.1Q	VLAN ID	Name	Status	Voice	Jumbo
1		DEFAULT_VLAN	Port-based	No	Yes
5		VLAN5	Port-based	No	No
22		VLAN22	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo packets.

Figure 14-8. Example Listing of Static VLANs To Show Jumbo Status Per VLAN

Syntax: show vlans ports < port-list >

*Lists the static VLANs to which the specified port(s) belong, including the **Jumbo** column to indicate which VLANs are configured to support jumbo traffic. Entering only one port in < port-list > results in a list of all VLANs to which that port belongs. Entering multiple ports in < port-list > results in a superset list that includes the VLAN memberships of all ports in the list, even though the individual ports in the list may belong to different subsets of the complete VLAN listing. For example, if port 1 belongs to VLAN 1, port 2 belongs to VLAN 10, and port 3 belongs to VLAN 15, then executing this command with a < port-list > of **1-3** results in a listing of all three VLANs, even though none of the ports belong to all three VLANs. (Refer to figure 14-9.)*

ProCurve# show vlans ports 1-3

Status and Counters - VLAN Information - for ports 1-3

802.1Q	VLAN ID	Name	Status	Voice	Jumbo
1		DEFAULT_VLAN	Port-based	No	Yes
10		VLAN10	Port-based	No	No
15		VLAN15	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo packets.

Figure 14-9. Example of Listing the VLAN Memberships for a Range of Ports

Syntax: show vlans < vid >

This command shows port membership and jumbo configuration for the specified < vid >.

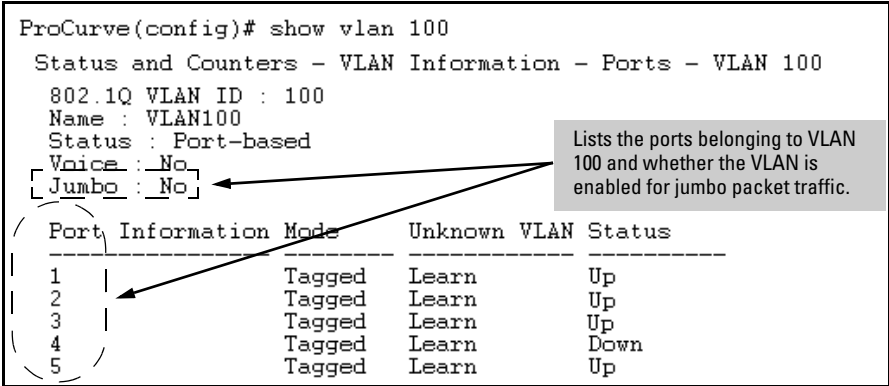


Figure 14-10. Example of Listing the Port Membership and Jumbo Status for a VLAN

Enabling or Disabling Jumbo Traffic on a VLAN

Syntax: vlan < vid > jumbo
[no] vlan < vid > jumbo

*Configures the specified VLAN to allow jumbo packets on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, **vlan < vid > jumbo** also creates the VLAN. Note that a port belonging to one jumbo VLAN can receive jumbo packets through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo packets. The **[no]** form of the command disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are **jumbo** and **no jumbo**. (Default: Jumbos disabled on the specified VLAN.)*

Operating Notes for Jumbo Traffic-Handling

- ProCurve does not recommend configuring a voice VLAN to accept jumbo packets. Voice VLAN packets are typically small, and allowing a voice VLAN to accept jumbo packet traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo packets on all ports belonging to the VLAN.
- When the switch applies the default MTU (1522-bytes) to a VLAN, all ports in the VLAN can receive incoming packets of up to 1522 bytes in length. When the switch applies the jumbo MTU (9220 bytes) to a VLAN, all ports in that VLAN can receive incoming packets of up to 9220 bytes in length. A port receiving packets exceeding the applicable MTU drops such packets, causing the switch to generate an Event Log message and increment the “Giant Rx” counter (displayed by **show interfaces < port-list >**).
- The switch does not allow flow control and jumbo packet capability to co-exist on a port. Attempting to configure both on the same port generates an error message in the CLI and sends a similar message to the Event Log.
- The default MTU on the 3400cl/6400cl switches is 1522 bytes (including 4 bytes for the VLAN tag). The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag).
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving “excessive” inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition generates a Fault-Finder message in the Alert log of the switch’s web browser interface, and also increments the switch’s “Giant Rx” counter.
- If you do not want all ports in a given VLAN to accept jumbo packets, you can consider creating one or more jumbo VLANs with a membership comprised of only the ports you want to receive jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo packets through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN. For example, suppose you wanted to allow inbound jumbo packets only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200, and also share these VLANs with other ports you want

excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-Enabled?	No	No	Yes

If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.

- **Outbound Jumbo Traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo packets through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo enabled VLANs. This can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo packets can forward them to the ports in the VLAN that do not have jumbo capability.

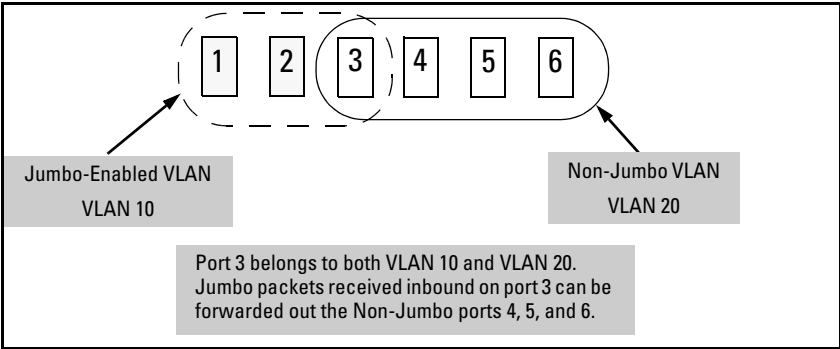


Figure 14-11. Forwarding Jumbo Packets Through Non-Jumbo Ports

Jumbo packets can also be forwarded out non-jumbo ports when the jumbo packets received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

- **Jumbo Traffic in a Switch Mesh Domain.** Note that if a switch belongs to a meshed domain, but does not have any VLANs configured to support jumbo traffic, then the meshed ports on that switch will drop any jumbo packets they receive from other devices. In this regard, if a mesh domain includes any ProCurve Series 5300xl switches and/or ProCurve 1600M/2400M/2424M/4000M/8000M switches along with Series 3400cl and Series 6400cl switches configured to support jumbo traffic, only the 3400cl/6400cl switches will receive jumbo packets. The other switch models in the mesh will drop such packets. For more information on switch meshing, refer to the chapter titled “Switch Meshing” in the *Advanced Traffic Management Guide* for your switch.

Troubleshooting

A VLAN is configured to allow jumbo packets, but one or more ports drops all inbound jumbo packets. The port may not be operating at 1 gigabit or higher. Regardless of a port’s configuration, if it is actually operating at a speed lower than 1 gigabit, it drops inbound jumbo packets. For example, if a port is configured for **Auto** mode (**speed-duplex auto**), but has negotiated a 100 Mbps speed with the device at the other end of the link, then the port cannot receive inbound jumbo packets. To determine the actual operating speed of one or more ports, view the **Mode** field in the output for the following command:

show interfaces brief <port-list>

A non-jumbo port is generating “Excessive undersize/giant packets” messages in the Event Log. The 3400cl/6400cl switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo packets received on the jumbo VLAN to non-jumbo ports. Refer to “Outbound Jumbo Traffic” on page 14-33.

Configuring for Network Management Applications

Contents

Using SNMP Tools To Manage the Switch	15-3
Overview	15-3
SNMP Management Features	15-4
Configuring for SNMP Access to the Switch	15-4
Configuring for SNMP Version 3 Access to the Switch	15-5
SNMP Version 3 Commands	15-6
Enabling SNMPv3	15-7
SNMPv3 Users	15-7
Group Access Levels	15-10
SNMPv3 Communities	15-11
Menu: Viewing and Configuring non-SNMP version 3 Communities	15-13
CLI: Viewing and Configuring SNMP Community Names	15-14
SNMPv3 Notification and Traps	15-16
SNMPv1 and SNMPv2c Trap Features	15-19
CLI: Configuring and Displaying Trap Receivers	15-19
Using the CLI To Enable Authentication Traps	15-22
Advanced Management: RMON	15-22
LLDP (Link-Layer Discovery Protocol)	15-24
Terminology	15-25
General LLDP Operation	15-27
LLDP-MED	15-27
Packet Boundaries in a Network Topology	15-27
Configuration Options	15-28
Options for Reading LLDP Information Collected by the Switch ..	15-30
LLDP and LLDP-MED Standards Compatibility	15-31

LLDP Operating Rules	15-31
LLDP Data Management on the Series 3400cl and 6400cl Switches	15-32
LLDP Neighbor Data	15-32
Configuring LLDP Operation	15-33
Viewing the Current Configuration	15-34
Configuring Global LLDP Packet Controls	15-37
Configuring SNMP Notification Support	15-41
Changing the Minimum Interval for Successive Data Change Notifications for the Same Neighbor	15-42
Configuring Per-Port Transmit and Receive Modes	15-42
Configuring Basic LLDP Per-Port Advertisement Content	15-43
Configuring Support for Port Speed and Duplex Advertisements on the 5300xl and 4200vl Switches	15-46
LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches	15-47
LLDP-MED Topology Change Notification	15-50
LLDP-MED Fast Start Control	15-51
Advertising Device Capability, Network Policy, PoE Status and Location Data	15-52
Configuring Location Data for LLDP-MED Devices	15-56
Displaying Advertisement Data	15-62
Displaying Switch Information Available for Outbound Advertisements	15-63
Displaying LLDP Statistics	15-68
LLDP Operating Notes	15-70
LLDP and CDP Data Management	15-72
LLDP and CDP Neighbor Data	15-72
CDP Operation and Commands	15-74

Using SNMP Tools To Manage the Switch

Overview

You can manage the switch via SNMP from a network management station running an application such as ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+). For more on PCM and PCM+, visit the ProCurve Networking web site at:

www.procurve.com

Click on **products index** in the sidebar, then click on the appropriate link appearing under the **Network Management** heading.

This section includes:

- An overview of SNMP management for the switch
- Configuring the switches for:
 - SNMP Communities (page 15-11)
 - Trap Receivers and Authentication Traps (page 15-16)
- Information on advanced management through RMON Support (page 15-22)

To implement SNMP management, the switch must have an IP address, configured either manually or dynamically (using DHCP or Bootp). If multiple VLANs are configured, each VLAN interface should have its own IP address. For DHCP use with multiple VLANs, refer to the section titled “The Primary VLAN” in the “Static Virtual LANs (VLANs)” chapter of the Advanced Traffic Management Guide for your switch.

Note

If you use the switch’s Authorized IP Managers and Management VLAN features, ensure that the SNMP management station and/or the choice of switch port used for SNMP access to the switch are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked. For more on Authorized IP Managers, refer to the Access Security Guide on the Documentation CD-ROM shipped with your switch and also available on the ProCurve Networking web site. For information on the Management VLAN feature, refer to the section titled “The Secure Management VLAN” in the “Static Virtual LANs (VLANs)” chapter of the *Advanced Traffic Management Guide* for your switch.

SNMP Management Features

SNMP management features on the switch include:

- SNMP version 1, version 2c or version 3 over IP
- Security via configuration of SNMP communities (page 15-4)
- Security via authentication and privacy for SNMP Version 3 access
- Event reporting via SNMP
 - Version 1 traps
 - RMON: groups 1, 2, 3, and 9
- ProCurve Manager/Plus support
- Flow sampling using either EASE or sFlow
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in a Hewlett-Packard proprietary MIB (Management Information Base) file. If you are using HP OpenView, you can ensure that it is using the latest version of the MIB file by downloading the file to the OpenView database. To do so, go to the ProCurve Networking web site at:

www.procurve.com

Click on **software updates**, then **MIBs**.

Configuring for SNMP Access to the Switch

SNMP access requires an IP address and subnet mask configured on the switch. (See “IP Configuration” on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See “DHCP/Bootp Operation” on page 8-12.)

Once an IP address has been configured, the main steps for configuring SNMP version 1 and version 2c access management features are:

1. Configure the appropriate SNMP communities. (Refer to “SNMPv3 Communities” on page 15-11.)
2. Configure the appropriate trap receivers. (Refer to “SNMPv3 Notification and Traps” on page 15-16.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community.

If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (Refer to the *Access Security Guide* for your switch.)

Caution

For ProCurve Manager (PCM) version 1.5 or earlier (or any TopTools version), deleting the “public” community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, ProCurve recommends that you change the write access for the “public” community to “Restricted”.

Configuring for SNMP Version 3 Access to the Switch

SNMP version 3 (SNMPv3) access requires an IP address and subnet mask configured on the switch. (See “IP Configuration” on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See “DHCP/Bootp Operation” on page 8-12.)

Once an IP address has been configured, the main steps for configuring SNMP version 3 access management features are:

1. Enable SNMPv3 for operation on the switch (Refer to “SNMP Version 3 Commands” on page 15-6)
2. Configure the appropriate SNMP users (Refer to “SNMPv3 Users” on page 15-7)
3. Configure the appropriate SNMP communities. (Refer to “SNMPv3 Communities” on page 15-11.)
4. Configure the appropriate trap receivers. (Refer to “SNMPv3 Notification and Traps” on page 15-16.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct User and community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (Refer to the *Access Security Guide* for your switch.)

SNMP Version 3 Commands

SNMP version 3 (SNMPv3) adds a new command to the CLI for configuring SNMPv3 functions. To enable SMNPv3 operation on the switch you must:

- a. Enable SNMPv3 with the **snmpv3 enable** command. An initial user entry will be generated with MD5 authentication and DES privacy.
- b. You may restrict access to only SNMPv3 agents with the **snmpv3 only** command. A second option would be to restrict write access to only SNMPv3 agents with the **snmpv3 restricted-access** command

Caution

Restricting access to only version 3 messages will make the community named “public” inaccessible to network management applications (such as auto-discovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

Syntax: [no] snmpv3 enable

Enable and disable the switch for access from SNMPv3 agents. This includes the creation of the initial user record.

[no] snmpv3 only

Enables or disables restrictions to access from only SNMPv3 agents. When enabled, the switch will reject all non-SNMPv3 messages.

[no] snmpv3 restricted-access

Enables or disables restrictions from all non-SNMPv3 agents to read only access.

show snmpv3 enable

Displays the operating status of SNMPv3.

show snmpv3 only

Displays status of message reception of non-SNMPv3 messages.

show snmpv3 restricted-access

Displays status of write messages of non-SNMPv3 messages.

Enabling SNMPv3

The **snmpv3 enable** command starts a dialog that performs three functions: enabling the switch to receive SNMPv3 messages, configuring the initial users, and, optionally, to restrict non-version 3 messages to “read only”. Figure 15-1 shows an example of this dialog.

Note:
SNMP
Version 3
Initial Users

For most SNMPv3 management software to be able to create new users, they must have an initial user record clone. These records can be downgraded, given less features, but not upgraded with new features added. For this reason it is recommended that a second user with SHA and DES are created at the time you enable SNMPv3

```
ProCurve (config)# snmpv3 enable
SNMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User 'initial' is created
Would you like to create a user that uses SHA? y
Enter user name: templateSHA
Authentication Protocol: SHA
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User creation is done.  SNMPv3 is now functional.
Would you like to restrict SNMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmp restrict-access'): n
```

Figure 15-1. Example of SNMP version 3 Enable Command

SNMPv3 Users

The second step to use SNMPv3 on the switch is to configure the users that will be assigned to different groups. To establish users on the switch:

1. Add the users to the User Table. This is done with the **snmpv3 user** command. To view the users in the list you use the **show snmpv3 user** command. See “Adding Users” on page 15-8.

2. Assign users to Security Groups based on their security model. This is done with the **snmpv3 group** command. See “Assigning Users to Groups” on page 15-9.

Caution

Adding a user without authentication and/or privacy to a group that requires it, will cause the user to not be able to access the switch. You should only add users to the group that is appropriate for their security parameters

Adding Users. To establish a user you must first add the user names to the list of known users. Add user names with the **snmpv3 user** CLI command.

The diagram illustrates the configuration of SNMPv3 users on an HP Switch. It shows the following CLI commands and their corresponding actions:

- `HP Switch (config)# snmpv3 user NetworkAdmin`: Add user NetworkAdmin with no Authentication or Privacy.
- `HP Switch (config)# [snmpv3 user NetworkMgr auth md5 authpass priv privpass]`: Add user NetworkMgr with authentication and privacy. The command is shown with brackets and underscores to indicate the structure.
- `HP Switch (config)# show snmpv3 user`: Display the configuration for all SNMPv3 users.

The output of the `show snmpv3 user` command is as follows:

```
Status and Counters - SNMP v3 Global Configuration Information
```

User Name	Auth. Protocol	Privacy Protocol
NetworkAdmin	None	None
NetworkMgr	MD5	des
initial	MD5	des
templateSHA	SHA	des

Figure 15-2. Adding and showing Users for SNMPv3

SNMPv3 Commands

Syntax: [no] snmpv3 user <user_name>

Adds or Deletes a user entry for snmpv3. Authorization and privacy are optional, but to use privacy, you must use authorization. When deleting a user, only the **user_name** is required.

[auth <md5 | sha> <auth_pass>]

With authorization, you can select either MD5 authentication or sha authentication. The **auth_pass** must be 6-32 characters in length and must be included when authentication is included. (Default: None)

[priv <priv_pass>]

With privacy, the switch only supports DES (56-bit) encryption. The privacy password **priv_pass** must be 6-32 characters in length and must be included when using the **priv** parameter. (Default: None)

Assigning Users to Groups. Then you must set the group access level for the user by assigning the user to a group. This is done with the **snmpv3 group** command. For more details on the MIBs access for a given group see “Group Access Levels” on page 15-10.

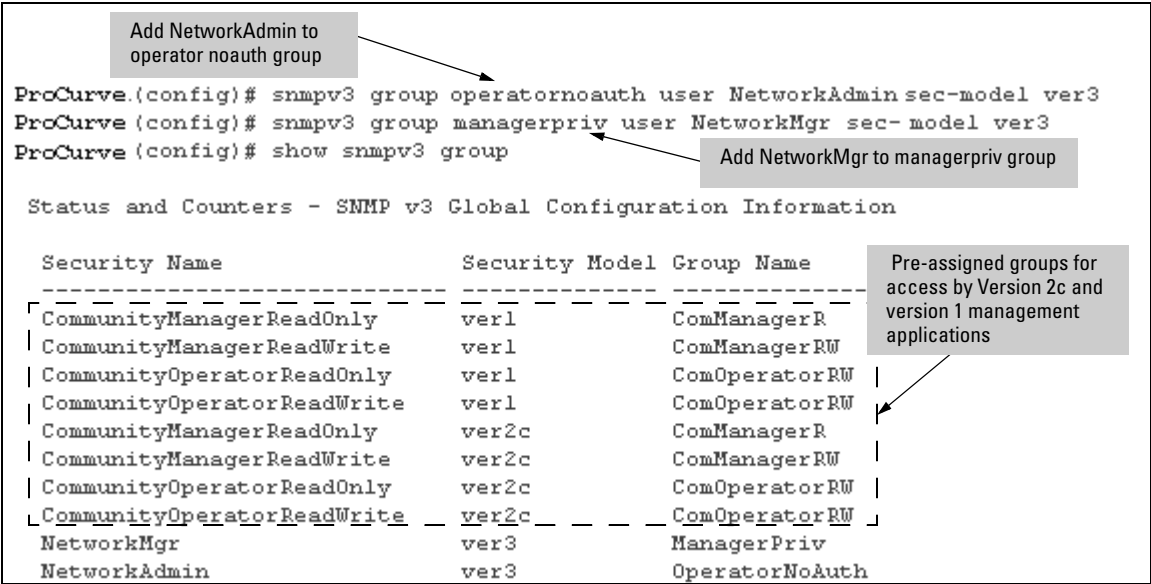


Figure 15-3. Example of Assigning Users to Groups

SNMPv3 Group Commands

Syntax: [no] snmpv3 group

This command assigns or removes a user to a security group for access rights to the switch. To delete an entry, all of the following three parameters must be included in the command.

group <group_name>

This parameter identifies the group that has the privileges that will be assigned to the user. For more details see “Group Access Levels” on page 15-10.

user <user_name>

*This parameter identifies the user to be added to the access group. This must match the user name added with the **snmpv3 user** command.*

sec-model <ver1 | ver2c | ver3>

This defines which security model to use for the added user. A SNMPv3 access Group should only use the ver3 security model.

Group Access Levels

The switch supports eight predefined group access levels. There are four levels for use with version 3 users and four are used for access by version 2c or version 1 management applications.

Group Name	Group Access Type	Group Read View	Group Write View
managerpriv	Ver3 Must have Authentication and Privacy	ManagerReadView	ManagerWriteView
managerauth	Ver3 Must have Authentication	ManagerReadView	ManagerWriteView
operatorauth	Ver3 Must have Authentication	OperatorReadView	DiscoveryView
operatornoauth	Ver3 No Authentication	OperatorReadView	DiscoveryView
commanagerrw	Ver2c or Ver1	ManagerReadView	ManagerWriteView
commanagerr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorr	Ver2c or Ver1	OperatorReadView	DiscoveryView

Each view allows you to view or modify a different set of MIBs.

- **Manager Read View** – access to all managed objects
- **Manager Write View** – access to all managed objects *except* the following: vacmContextTable, vacmAccessTable, vacmViewTreeFamilyTable

- **OperatorReadView** – no access to icfSecurityMIB, hpSwitchIpTftp-Mode, vacmContextTable, vacmAccessTable, vacmViewTreeFamilyTable, usmUserTable, snmpCommunityTable
- **Discovery View** – Access limited to samplingProbe MIB.

Note

All access groups and views are predefined on the switch. There is no method to modify or add groups or views to those that are pre-defined on the switch.

SNMPv3 Communities

SNMP communities are supported by the switch to allow management application that use version 2c or version 1 to access the switch. The communities are mapped to Group Access Levels that are used for version 2c or version 1 support. For more information see “Group Access Levels” on page 15-10. This mapping will happen automatically based on the communities access privileges, but special mappings can be added with the **snmpv3 community** command.

Syntax: [no] snmpv3 community

*This command maps or removes a mapping of a community name to a group access level. To remove a mapping you, only need to specify the **index_name** parameter.*

index <index_name>

This is an index number or title for the mapping. The values of 1-5 are reserved and can not be mapped.

name <community_name>

This is the community name that is being mapped to a group access level.

sec-name <security_name>

This is the group level that the community is being mapped. For more information see “Group Access Levels” on page 15-10.

tag <tag_value>

This is used to specify which target address may have access by way of this index reference.

Figure 15-4 shows the assigning of Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator only has an access level of CommunityOperatorReadOnly

```
ProCurve (config)# snmpv3 community index 30 name Operator sec-name
CommunityManagerReadWrite tag MgrStation1
ProCurve (config)# show snmpv3 community
```

Index	Name	Community Name	Security Name
1	public	CommunityManagerReadWrite	CommunityManagerReadWrite
2	Operator	CommunityOperatorReadOnly	CommunityOperatorReadOnly
3	Manager	CommunityManagerReadWrite	CommunityManagerReadWrite
30	Operator	CommunityManagerReadWrite	CommunityManagerReadWrite

Figure 15-4. Assigning a Community to a Group Access Level

SNMP Community Features

Feature	Default	Menu	CLI	Web
show SNMP communities	n/a	page 15-13	page 15-14	—
configure identity information	none	—	page 15-15	—
configure community names	public	page 15-13	page 15-15	—
MIB view for a community name (operator, manager)	manager	"	"	—
write access for default community name	unrestricted	"	"	—

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view, and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.

Caution

For ProCurve Manager (PCM) version 1.5 or earlier (or any TopTools version), deleting the “public” community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, ProCurve recommends that you change the write access for the “public” community to “Restricted”.

**Menu: Viewing and Configuring non-SNMP
version 3 Communities**

To View, Edit, or Add SNMP Communities:

1. From the Main Menu, Select:
2. Switch Configuration...
6. SNMP Community Names

Note: This screen gives an overview of the SNMP communities that are currently configured. All fields in this screen are read-only.

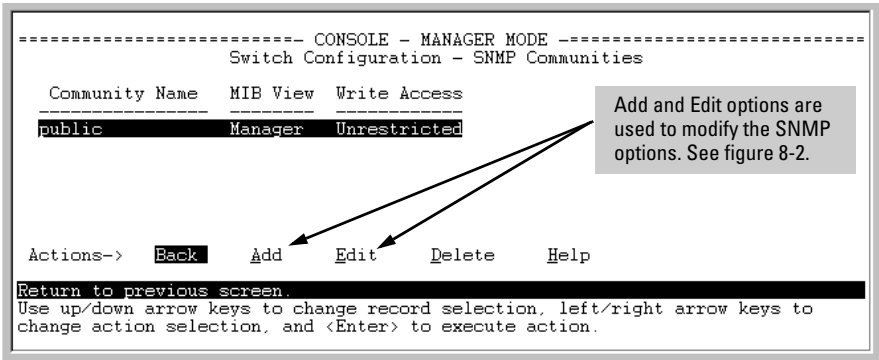


Figure 15-5. The SNMP Communities Screen (Default Values)

2. Press [A] (for **Add**) to display the following screen:

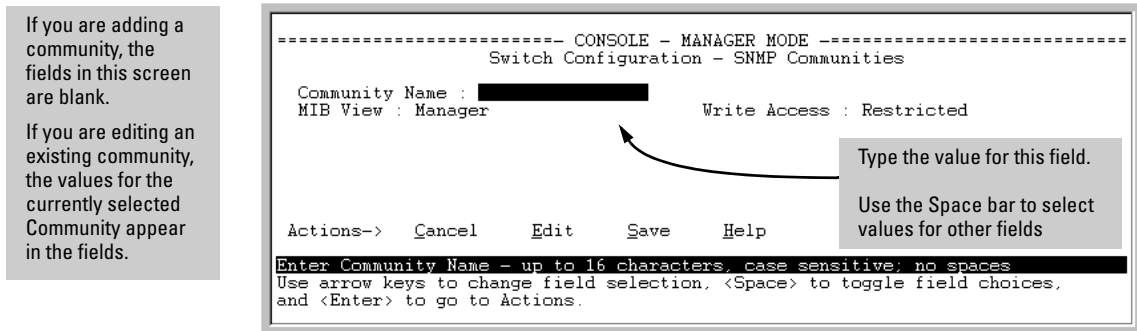


Figure 15-6. The SNMP Add or Edit Screen

Need Help? If you need information on the options in each field, press **[Enter]** to move the cursor to the Actions line, then select the **H**elp option on the Actions line. When you are finished with Help, press **[E]** (for **E**dit) to return the cursor to the parameter fields.

3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the **[Tab]** key to move from one field to the next.)
4. Press **[Enter]**, then **[S]** (for **S**ave).

CLI: Viewing and Configuring SNMP Community Names

Community Name Commands	Page
show snmp-server [<community-string>]	15-14
[no] snmp-server	15-15
[community <community-str>]	15-15
[host <community-str> <ip-addr>]	15-20
[<none debug all not-info critical>]	
[enable traps <authentication>]	15-22

Listing Community Names and Values. This command lists the data for currently configured SNMP community names (along with trap receivers and the setting for authentication traps — see “SNMPv3 Notification and Traps” on page 15-16).

Syntax: show snmp-server [<community-string>]

This example lists the data for all communities in a switch; that is, both the default “public” community name and another community named “blue-team”

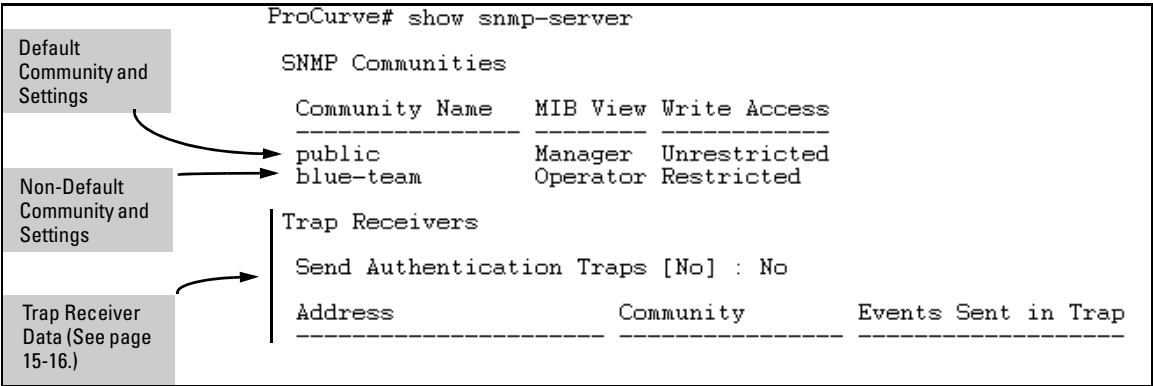


Figure 15-7. Example of the SNMP Community Listing with Two Communities

To list the data for only one community, such as the "public" community, use the above command with the community name included. For example:

```
ProCurve# show snmp-server public
```

Configuring Community Names and Values. The **snmp-server** command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax: [no] snmp-server community < community-name >

*Configures a new community name. If you do not also specify **operator** or **manager**, the switch automatically assigns the community to the **operator** MIB view. If you do not specify **restricted** or **unrestricted**, the switch automatically assigns the community to **restricted** (read-only) access. The **no** form uses only the < community-name > variable and deletes the named community from the switch.*

[operator | manager]

*Optionally assigns an access level. At the **operator** level the community can access all MIB objects except the CONFIG MIB. At the **manager** level the community can access all MIB objects.*

[restricted | unrestricted]

*Optionally assigns MIB access type. Assigning the **restricted** type allows the community to read MIB variables, but not to set them. Assigning the **unrestricted** type allows the community to read and set MIB variables.*

For example, to add the following communities:

Community	Access Level	Type of Access
red-team	manager (Access to all MIB objects.)	unrestricted (read/write)
blue-team	operator (Access to all MIB objects except the CONFIG MIB.)	restricted (read-only)

```
ProCurve(config)# snmp-server community red-team  
manager unrestricted
```

```
ProCurve(config)# snmp-server community blue-team  
operator restricted
```

To eliminate a previously configured community named "gold-team":

```
ProCurve(config) # no snmp-server community gold-team
```

SNMPv3 Notification and Traps

The switches covered by this manual support the SNMPv3 notification process. They also support version 1 or version 2c traps. For more information on version 1 or version 2c traps, see "SNMPv1 and SNMPv2c Trap Features" on page 15-19. The SNMPv3 notification process allows for the messages passed to be authenticated and encrypted if you choose. To set up a SNMPv3 notification there are three steps:

1. Establish a Notification with the **snmpv3 notify** command
2. Point the notification to an Address with the **snmpv3 targetaddress** command.
3. Establish a parameter record for the target address with the **snmpv3 params** command.

Syntax: [no] snmpv3 notify <notify_name> tag <tag_name>

*This adds or deletes a notification request. To remove a mapping you only need the **< notify_name >**.*

[no] snmpv3 targetaddress < name > taglist < tag > params
< parms_name > < ip-addr >

*Add or delete an address where notification messages are sent. The **< tag >** value must match the tag value of a notify entry.*

Note: You are only allowed up to 103 characters for the **taglist** value.

filter < none | debug | all | not-info | critical>

*This filters messages to restrict the types of messages transmitted to an address. (Default: **none**)*

udp-port < port >

*This specifies the UDP port to use. (Default: **162**)*

port-mask < mask >

*Used to specific a range of UDP ports. (Default: **0**)*

addr-mask < mask >

*Used to specify a range of addresses as destinations for notify messages. (Default: **0**)*

retries < value >

*Number of times to retransmit a message when no response is reviewed. (Default: **3**)*

timeout < value >

*Specifies how long the switch waits for a response from the target before it retransmits the packet. (Default: **1500**)*

max-msg-size<size> **Default:1472**

Specifies the maximum number of bytes a message to this target can contain.

[no] snmpv3 params <params_name> user <user_name>

*Adds or deletes a user parameter for use with target address. The **params_name** must match the **parms_name** in the **targetaddress** command. The **user_name** should be a user from the User Table. For more information on users see “SNMPv3 Users” on page 15-7.*

*A complete **params** command must also have a **sec-model** and **msg-processing** entry.*

< sec-model < ver1 | ver2c | ver3 >

*This established the security model to use for messages passed to the **targetaddress**. If you use **ver3** then **msg-processing** must also be **ver3**.*

< msg-processing < ver1 | ver2c | ver3 > [noaut | auth | priv]

*Establishes the **msg-processing** algorithm for messages passed to the target address. If **ver3** is used and **sec-model** is **ver3** then you must select a security services level (**noauth**, **auth**, or **priv**).*

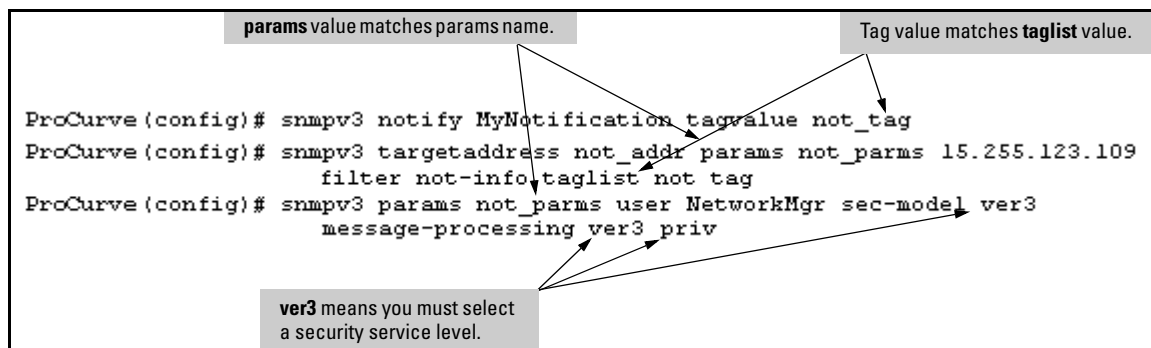


Figure 15-8. Example of SNMP Notification and Trap Configuration

SNMPv1 and SNMPv2c Trap Features

Feature	Default	Menu	CLI	Web
snmp-server host (trap receiver)	public	—	page 15-20	—
snmp-server enable (authentication trap)	none	—	page 15-22	—

A *trap receiver* is a management station designated by the switch to receive SNMP traps sent from the switch. An *authentication trap* is a specialized SNMP trap sent to trap receivers when an unauthorized management station tries to access the switch.

Note

Fixed or “Well-Known” Traps: The Series 5300xl switches automatically sends fixed traps (such as “coldStart”, “warmStart”, “linkDown”, and “linkUp”) to trap receivers using the **public** community name. These traps cannot be redirected to other communities. Thus, if you change or delete the default **public** community name, these traps will be lost.

Thresholds: The switch automatically sends all messages resulting from thresholds to the network management station(s) that set the thresholds, regardless of the trap receiver configuration.

In the default configuration, there are no trap receivers configured, and the authentication trap feature is disabled. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch. As an option, you can also configure the switch to send Event Log messages as traps.

CLI: Configuring and Displaying Trap Receivers

Trap Receiver Commands	Page
show snmp-server	15-20
snmp-server host <ip-addr> <community-name> [none all non-inform critical debug]	15-20
snmp-server enable traps authentication	15-20

Using the CLI To List Current SNMP Trap Receivers.

This command lists the currently configured trap receivers and the setting for authentication traps (along with the current SNMP community name data — see “SNMPv3 Communities” on page 15-11).

Syntax: show snmp-server

Displays current community and trap receiver data.

In the next example, the **show snmp-server** command shows that the switch has been previously configured to send SNMP traps to management stations belonging to the “public”, “red-team”, and “blue-team” communities.

```
ProCurve# show snmp-server
```

SNMP Communities			
Community Name	MIB View	Write Access	
public	Operator	Restricted	
blue-team	Manager	Unrestricted	
red-team	Manager	Unrestricted	

Trap Receivers			
Send Authentication Traps : No			
Address	Community	Events Sent in Trap	
10.28.227.200	public	All	
10.28.227.105	red-team	Critical	
10.28.227.120	blue-team	Not-INFO	

Figure 15-9. Example of Show SNMP-Server Listing

Configuring Trap Receivers. This command specifies trap receivers by community membership, management station IP address, and the type of Event Log messages to send to the trap receiver.

Note

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps will be sent to that trap receiver until the community to which it belongs has been configured on the switch.

Syntax: snmp-server host < community-string > < ip-address >

Using community name and destination IP address, this command designates a destination network-management station for receiving SNMP event log messages from the switch. If you do not specify the event level, then the switch does not send event log messages as traps. You can specify up to 10 trap receivers (network management stations).

Note: In all cases, the switch sends any threshold trap(s) to the network management station(s) that explicitly set the threshold(s).

[<none | all | non-info | critical | debug>]

Options for sending switch Event Log messages to a trap receiver. Refer to Table 15-1, “Options for Sending Event Log Messages as Traps,” on page 15-21. The levels specified with these options apply only to Event Log messages, and not to threshold traps.

Table 15-1. Options for Sending Event Log Messages as Traps

Event Level	Description
None (default)	Send no log messages.
All	Send all log messages.
Not INFO	Send the log messages that are not information-only.
Critical	Send critical-level log messages.
Debug	Reserved for ProCurve-internal use.

For example, to configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" log messages:

```
ProCurve(config)# snmp-server trap-receiver red-team
10.28.227.130 critical
```

Notes

To replace one community name with another for the same IP address, you must use **no snmp-server host < community-name> < ip-address >** to delete the unwanted community name. Otherwise, adding a new community name with an IP address already in use with another community name simply creates two allowable community name entries for the same management station.

If you do not specify the event level ([<none | all | non-info | critical | debug>]) then the switch does not send event log messages as traps. "Well-Known" traps and threshold traps (if configured) will still be sent.

Using the CLI To Enable Authentication Traps

Note

For this feature to operate, one or more trap receivers must be configured on the switch. See "Configuring Trap Receivers" on page 15-20.

Using the CLI To Enable Authentication Traps.

Syntax: [no] snmp-server enable traps authentication

Enables or disables sending an authentication trap to the configured trap receiver(s) if an unauthorized management station attempts to access the switch.

For example:

```
ProCurve(config)# snmp-server enable traps authentication
```

Check the Event Log in the console interface to help determine why the authentication trap was sent. (Refer to "Using the Event Log To Identify Problem Sources" on page C-27.)

Advanced Management: RMON

The switch supports RMON (Remote Monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the ProCurve Manager network management software. For more on ProCurve Manager, visit the ProCurve Networking web site at

www.procurve.com

Click on **products index**, then look for the ProCurve Manager topic under the **Network Manager** bar.

LLDP (Link-Layer Discovery Protocol)

To standardize device discovery on all ProCurve switches, LLDP will be implemented while offering limited read-only support for CDP as documented in this manual. For current information on your switch model, consult the Release Notes (available on the ProCurve Networking web site). If LLDP has not yet been implemented (or if you are running an older version of software), consult a previous version of the Management and Configuration Guide for device discovery details.

Note

LLDP-Med features are supported on the Series 5300xl and 4200vl switches.

Table 15-2. LLDP and LLDP-MED Features

Feature	Default	Menu	CLI	Web
View the switch's LLDP configuration	n/a	—	page 15-33	—
Enable or disable LLDP on the switch	Enabled	—	page 15-28	—
Change the transmit interval (refresh-interval) for LLDP packets	30 seconds	—	page 15-38	—
Change the holdtime multiplier for LLDP Packets (holdtime-multiplier x refresh-interval = time-to-live)	4 seconds	—	page 15-28	—
Change the delay interval between advertisements	2 seconds	—	page 15-39	—
Changing the reinitialization delay interval	2 seconds	—	page 15-40	—
Configuring SNMP notification support	Disabled	—	page 15-41	—
Configuring transmit and receive modes	tx_rx	—	page 15-42	—
Configuring basic LLDP per-port advertisement content	Enabled	—	page 15-43	—
Configuring port speed and duplex advertisements for optional LLDP and mandatory LLDP-MED applications	Enabled	—	page 15-65	—
Configuring topology change notification for LLDP-MED	Enable	—	page 15-71	—
Changing the fast-start duration for LLDP-MED	5 sec	—	page 15-51	—
Configuring LLDP-MED Advertising	Enabled	—	page 15-43	—
Configuring LLDP-MED device location data	None	—	page 15-63	—
Displaying Advertisement Data and Statistics	n/a	—	page 15-68	—

LLDP (Link Layer Discovery Protocol): provides a standards-based method for enabling the switches covered by this guide to advertise themselves to adjacent devices and to learn about adjacent LLDP devices.

LLDP-MED (LLDP Media Endpoint Discovery): Provides an extension to LLDP and is designed to support VoIP deployments.

Note

LLDP-MED is an extension for LLDP, and the switch requires that LLDP be enabled as a prerequisite to LLDP-MED operation. As of October, 2005, LLDP-MED operates on 5300xl and 4200vl switches. This feature is not currently offered on the 3400cl and 6400cl switches.

An SNMP utility can progressively discover LLDP devices in a network by:

1. Reading a given device's Neighbors table (in the Management Information Base, or MIB) to learn about other, neighboring LLDP devices.
2. Using the information learned in step 1 to find and read the neighbor devices' Neighbors tables to learn about additional devices, and so on.

Also, by using **show** commands to access the switch's neighbor database for information collected by an individual switch, system administrators can learn about other devices connected to the switch, including device type (capability) and some configuration information. In VoIP deployments using LLDP-MED on the 5300xl or 4200vl switches, additional support unique to VoIP applications is also available. Refer to "LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches" on page 15-47.

Terminology

Adjacent Device: Refer to "Neighbor or Neighbor Device".

Advertisement: See LLDPDU.

Active Port: A port linked to another active device (regardless of whether STP is blocking the link).

ELIN (Emergency Location Identification Number): A valid telephone number in the North American Numbering Plan format and assigned to a multiline telephone system operator by the appropriate authority. This number calls a public service answering point (PSAP) and relays automatic location identification data to the PSAP.

LLDP: Link Layer Discovery Protocol:

- 5300xl, 4200vl, and 6400cl Switches: IEEE 802.1AB
- 3400cl Switches: IEEE 802.1AB/D9 or greater

LLDP-Aware: A device that has LLDP in its operating code, regardless of whether LLDP is enabled or disabled.

LLDP Device: A switch, server, router, or other device running LLDP.

LLDP Neighbor: An LLDP device that is either directly connected to another LLDP device or connected to that device by another, non-LLDP Layer 2 device (such as a hub) Note that an 802.1D-compliant switch does not forward LLDP data packets even if it is not LLDP-aware.

LLDPDU (LLDP Data Unit): LLDP data packet are transmitted on active links and include multiple TLVs containing global and per-port switch information. In this guide, LLDPDUs are termed “advertisements” or “packets”.

LLDP-MED (Link Layer Discover Protocol Media Endpoint Discovery): The TIA telecommunications standard produced by engineering subcommittee TR41.4, “VoIP Systems — IP Telephony infrastructure and Endpoints” to address needs related to deploying VoIP equipment in IEEE 802-based environments. This standard will be published as ANSI/TIA-1057.

MIB (Management Information Base): An internal database the switch maintains for configuration and performance information.

MLTS (Multiline Telephone System): A network-based and/or premises-based telephone system having a common interface with the public switched telephone system and having multiple telephone lines, common control units, multiple telephone sets, and control hardware and software.

NANP (North American Numbering Plan): A ten-digit telephone number format where the first three digits are an area code and the last seven-digits are a local telephone number.

Neighbor: See “LLDP Neighbor”.

Non-LLDP Device: A device that is not capable of LLDP operation.

PD (Powered Device): This is an IEEE 802.3af-compliant device that receives its power through a direct connection to a 10/100Base-TX PoE RJ-45 port in a ProCurve fixed-port or chassis-based switch. Examples of PDs include Voice-over-IP (VoIP) telephones, wireless access points, and remote video cameras.

PSAP (Public Safety Answering Point): PSAPs are typically emergency telephone facilities established as a first point to receive emergency (911) calls and to dispatch emergency response services such as police, fire and emergency medical services.

PSE (Power-Sourcing Equipment): A PSE, such as a PoE module installed in a Series 5300xl switch, provides power to IEEE 802.3af-compliant PDs directly connected to the ports on the module.

TLV (Type-Length-Value): A data unit that includes a data type field, a data unit length field (in bytes), and a field containing the actual data the unit is designed to carry (as an alphanumeric string, a bitmap, or a subgroup of information). Some TLVs include subelements that occur as separate data points in displays of information maintained by the switch for LLDP advertisements. (That is, some TLVs include multiple data points or subelements.)

General LLDP Operation

An LLDP packet contains data about the transmitting switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled, and reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. (LLDP is a one-way protocol and does not include any acknowledgement mechanism.) An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a Neighbor database (MIB).

LLDP-MED

This capability is an extension to LLDP and is available on 5300xl switches running software release E.10.02 or greater, and on 4200vl switches. Refer to “LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches” on page 15-47.

Packet Boundaries in a Network Topology

- Where multiple LLDP devices are directly connected, an outbound LLDP packet travels only to the next LLDP device. An LLDP-capable device does not forward LLDP packets to any other devices, regardless of whether they are LLDP-enabled.

- An intervening hub or repeater forwards the LLDP packets it receives in the same manner as any other multicast packets it receives. Thus, two LLDP switches joined by a hub or repeater handle LLDP traffic in the same way that they would if directly connected.
- Any intervening 802.1D device or Layer-3 device that is either LLDP-unaware or has disabled LLDP operation drops the packet.

Configuration Options

Enable or Disable LLDP on the Switch. In the default configuration, LLDP is globally enabled on the switch. To prevent transmission or receipt of LLDP traffic, you can disable LLDP operation (page 15-28)

Enable or Disable LLDP-MED on 5300xl or 4200vl Switches. In the default configuration for 5300xl switches running software release E.10.02 or greater or 4200vl switches, LLDP-MED is enabled by default. (Requires that LLDP is also enabled.) For more information, refer to “LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches” on page 15-47.

Change the Frequency of LLDP Packet Transmission to Neighbor Devices. On a global basis, you can increase or decrease the frequency of outbound LLDP advertisements (page 15-28).

Change the Time-To-Live for LLDP Packets Sent to Neighbors. On a global basis, you can increase or decrease the time that the information in an LLDP packet outbound from the switch will be maintained in a neighbor LLDP device (page 15-28).

Transmit and Receive Mode. With LLDP enabled, the switch periodically transmits an LLDP advertisement (packet) out each active port enabled for outbound LLDP transmissions, and receives LLDP advertisements on each active port enabled to receive LLDP traffic (page 15-42). Per-Port configuration options include four modes:

- Transmit and Receive (**tx_rx**): This is the default setting on all ports. It enables a given port to both transmit and receive LLDP packets, and to store the data from received (inbound) LLDP packets in the switch's MIB.
- Transmit only (**txonly**): This setting enables a port to transmit LLDP packets that can be read by LLDP neighbors. However, the port drops inbound LLDP packets from LLDP neighbors without reading them. This prevents the switch from learning about LLDP neighbors on that port.

- **Receive only (rxonly):** This setting enables a port to receive and read LLDP packets from LLDP neighbors, and to store the packet data in the switch's MIB. However, the port does not transmit outbound LLDP packets. This prevents LLDP neighbors from learning about the switch through that port.
- **Disable (disable):** This setting disables LLDP packet transmissions and reception on a port. In this state, the switch does not use the port for either learning about LLDP neighbors or informing LLDP neighbors of its presence.

SNMP Notification. You can enable the switch to send a notification to any configured SNMP trap receiver(s) when the switch detects a remote LLDP data change on an LLDP-enabled port (page 15-41).

Per-Port (Outbound) Data Options. The following table lists the information the switch can include in the per-port, outbound LLDP packets it generates. In the default configuration, all outbound LLDP packets include this information in the TLVs transmitted to neighbor devices. However, you can configure LLDP advertisements on a per-port basis to omit some of this information (page 15-43).

Table 15-3. Data Available for Basic LLDP Advertisements

Data Type	Configuration Options	Default	Description
Time-to-Live	See note 1.	120 Seconds	The length of time an LLDP neighbor retains the advertised data before discarding it.
Chassis Type ^{2, 6}	N/A	Always Enabled	Indicates the type of identifier used for Chassis ID.
Chassis ID ⁶	N/A	Always Enabled	Uses base MAC address of the switch.
Port Type ^{3, 6}	N/A	Always Enabled	Uses "Local", meaning assigned locally by LLDP.
Port Id ⁶	N/A	Always Enabled	Uses port number of the physical port. In the 5300xl switches, this is an internal number reflecting the reserved slot/port position in the chassis. For more information on this numbering scheme, refer to figures D-2 and D-3 in Appendix D, "MAC Address Management" of the <i>Management and Configuration Guide</i> for your switch.
Remote Management Address			
Type ^{4, 6}	N/A	Always Enabled	Shows the network address type.
Address ⁴	Default or Configured	Uses a default address selection method unless an optional address is configured. See "Remote Management Address", below.	
System Name ⁶	Enable/Disable	Enabled	Uses the switch's assigned name.

Data Type	Configuration Options	Default	Description
System Description ⁶	Enable/Disable	Enabled	Includes switch model name and running software version, and ROM version.
Port Description ⁶	Enable/Disable	Enabled	Uses the physical port identifier.
System capabilities supported ^{5, 6}	Enable/Disable	Enabled	Identifies the switch's primary capabilities (bridge, router).
System capabilities enabled ^{5, 6}	Enable/Disable	Enabled	Identifies the primary switch functions that are enabled, such as routing.

¹The Packet Time-to-Live value is included in LLDP data packets. (Refer to "Changing the Time-to-Live for Transmitted Advertisements" on page 15-38.)

²Subelement of the Chassis ID TLV.

³Subelement of the Port ID TLV.

⁴Subelement of the Remote-Management-Address TLV.

⁵Subelement of the System Capability TLV.

⁶Populated with data captured internally by the switch. For more on these data types, refer to the IEEE P802.1AB Standard.

Remote Management Address. The switch always includes an IP address in its LLDP advertisements. This can be either an address selected by a default process, or an address configured for inclusion in advertisements. Refer to "IP Address Advertisements" on page 15-31.

Debug Logging. You can enable LLDP debug logging to a configured debug destination (Syslog server and/or a terminal device) by executing the **debug lldp** command. (For more on Debug and Syslog, refer to the Troubleshooting appendix in the *Management and Configuration Guide* for your switch.) Note that the switch's Event Log does not record usual LLDP update messages.

Options for Reading LLDP Information Collected by the Switch

You can extract LLDP information from the switch to identify adjacent LLDP devices. Options include:

- Using the switch's **show lldp info** command options to display data collected on adjacent LLDP devices—as well as the local data the switch is transmitting to adjacent LLDP devices (page 15-33).

- Using an SNMP application that is designed to query the Neighbors MIB for LLDP data to use in device discovery and topology mapping. (In the 3400cl and 6400cl switches only.)
- Using the **walkmib** command to display a listing of the LLDP MIB objects

LLDP and LLDP-MED Standards Compatibility

The operation covered by this section is compatible with these standards:

- IEEE P802.1AB/D9 (Series 3400cl switches)
- IEEE P802.1AB (Series 5300xl, Series 4200vl, and Series 6400cl switches)
- RFC 2922 (PTOPO, or Physical Topology MIB)
- RFC 2737 (Entity MIB)
- RFC 2863 (Interfaces MIB)
- ANSI/TIA-1057/D6 (LLDP-MED; refer to “LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches” on page 15-47.)

LLDP Operating Rules

(For additional information specific to LLDP-MED operation, refer to “LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches” on page 15-47.)

Port Trunking. LLDP manages trunked ports individually. That is, trunked ports are configured individually for LLDP operation, in the same manner as non-trunked ports. Also, LLDP sends separate advertisements on each port in a trunk, and not on a per-trunk basis. Similarly, LLDP data received through trunked ports is stored individually, per-port.

IP Address Advertisements. In the default operation, if a port belongs to only one static VLAN, then the port advertises the lowest-order IP address configured on that VLAN. If a port belongs to multiple VLANs, then the port advertises the lowest-order IP address configured on the VLAN with the lowest VID. If the qualifying VLAN does not have an IP address, the port advertises 127.0.0.1 as its IP address. For example, if the port is a member of the default VLAN (VID = 1), and there is an IP address configured for the default VLAN, then the port advertises this IP address. In the default operation, the IP address that LLDP uses can be an address acquired by DHCP or Bootp.

You can override the default operation by configuring the port to advertise any IP address that is manually configured on the switch, even if the port does not belong to the VLAN configured with the selected IP address (page 15-43). (Note that LLDP cannot be configured through the CLI to advertise addresses acquired through DHCP or Bootp. However, as mentioned above, in the default LLDP configuration, if the lowest-order IP address on the VLAN with the lowest VID for a given port is a DHCP or Bootp address, then the switch includes this address in its LLDP advertisements unless another address is configured for advertisements on that port.) Also, although LLDP allows configuring multiple remote management addresses on a port, only the lowest-order address configured on the port will be included in outbound advertisements. Attempting to use the CLI to configure LLDP with an IP address that is either not configured on a VLAN, or has been acquired by DHCP or Bootp results in the following error message.

```
xxx.xxx.xxx.xxx: This IP address is not configured or is  
a DHCP address.
```

Spanning-Tree Blocking. Spanning tree does not prevent LLDP packet transmission or receipt on STP-blocked links.

802.1x Blocking. Ports blocked by 802.1x operation do not allow transmission or receipt of LLDP packets.

Data Management on the 3400cl/6400cl Switches. For information on how these switches manage LLDP data read from other devices, refer to “LLDP Data Management on the Series 3400cl and 6400cl Switches” on page 15-32.

LLDP Data Management on the Series 3400cl and 6400cl Switches

This section applies only to the Series 3400cl and 6400cl switches.

LLDP (Link-Layer Discovery Protocol) operation on the 3400cl and 6400cl switches includes transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices.) This section describes points to note regarding LLDP data received by a 3400cl or 6400cl switch from other devices.

LLDP Neighbor Data

With LLDP enabled on a switch port, the port can read LLDP advertisements, and stores the data from the advertisements in its neighbor database. If the switch receives LLDP advertisements on the same port from the same

neighbor, it stores this information as two separate entries if the advertisements have differences chassis ID and port ID information. However, if the chassis and port ID information are the same, the switch stores this information as a single entry.

LLDP data transmission/collection is enabled in the switch's default configuration. In this state, an SNMP network management application designed to discover devices running LLDP can retrieve neighbor information from the switch.

Table 15-4. 3400cl/6400cl Neighbor Data Management

Protocol State	Packet Generation	Inbound Data Management	Inbound Packet Forwarding
LLDP Enabled ¹	Generates and transmits LLDP packets out all ports on the switch.	Store inbound LLDP data.	No forwarding of inbound LLDP packets.
LLDP Disabled	No packet generation.	No storage of LLDP data from neighbor devices.	No forwarding of inbound LLDP packets.
¹ LLDP transmit/receive are enabled in the default configuration of 3400cl and 6400cl switches.			

Configuring LLDP Operation

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings that apply to all active ports on the switch, and per-port settings that affect only the operation of the specified ports.

The commands in this section affect both LLDP and LLDP-MED operation. for information on operation and configuration unique to LLDP-MED, refer to “LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches” on page 15-47.

Command	Page
show lldp config	15-36
[no] lldp run	15-37
lldp refresh-interval	15-38
lldp holdtime-multiplier	15-38
lldpTxDelay	15-39
lldpReinitDelay	15-40

Command	Page
lldp enable-notification	15-41
lldpnotificationinterval	15-42
lldp admin-status < txonly rxonly tx_rx disable >	15-42
lldp config < port-list > lPAddrEnable	15-43
lldp config < port-list > basicTlvEnable	15-44
lldp config < port-list > dot3TlvEnable < macphy_config >	15-46

Viewing the Current Configuration

Displaying the Global LLDP, Port Admin, and SNMP Notification Status. This command displays the switch's general LLDP configuration status, including some per-port information affecting advertisement traffic and trap notifications.

Syntax show lldp config

Displays the LLDP global configuration, LLDP port status, and SNMP notification status. For information on port admin status, refer to “Configuring Per-Port Transmit and Receive Modes” on page 15-42.

For example, **show lldp config** produces the following display when the switch is in the default LLDP configuration:

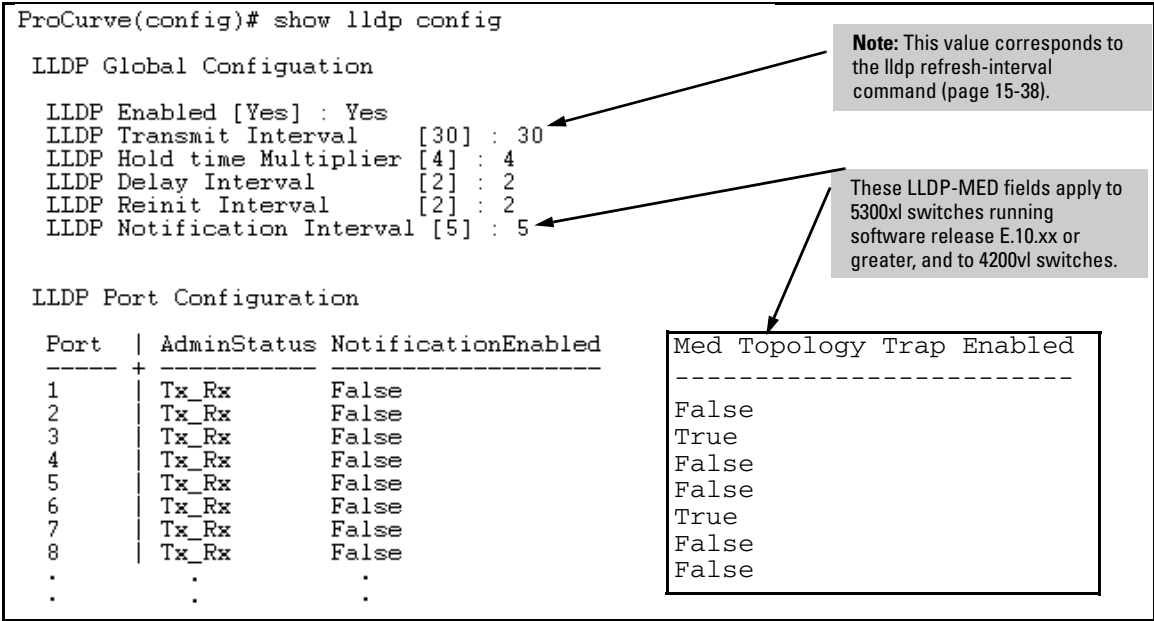


Figure 15-10. Example of Viewing the General LLDP Configuration

Displaying Port Configuration Details. This command displays the port-specific configuration.

Syntax `show lldp config < port-list >`

Displays the LLDP port-specific configuration for all ports in < port-list >, including which optional TLVs and any non-default IP address that are included in the port's outbound advertisements. For information on the notification setting, refer to "Configuring SNMP Notification Support" on page 15-41. For information on the other configurable settings displayed by this command, refer to "Configuring Per-Port Transmit and Receive Modes" on page 15-42.

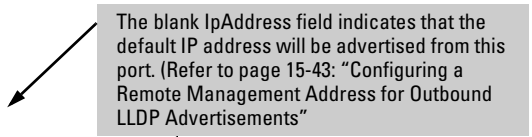
```
ProCurve(config)# show lldp config 3

LLDP Port Configuration Detail

Port : 3
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False

TLVS Advertised:
* port_descr
* system_name
* system_descr
* system_cap

IpAddress Advertised:
```



The blank IpAddress field indicates that the default IP address will be advertised from this port. (Refer to page 15-43: "Configuring a Remote Management Address for Outbound LLDP Advertisements")

Figure 15-11. Example of Per-Port Configuration Display (3400c1/6400c1 Switches)

```
ProCurve(config)# show lldp config a1

LLDP Port Configuration Detail

Port : a1
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

TLVS Advertised:
* port_descr
* system_name
* system_descr
* system_cap

[ * capabilities
  * network_policy
  * location_id
  * poe
  * macphy_config ]

IpAddress Advertised:
```

These fields appear when medtlvenable is enabled on the switch, which is the default setting.

This field appears when dot3tlvenable is enabled on the switch, which is the default setting.

The blank IpAddress field indicates that the default IP address will be advertised from this port. (Refer to page 15-43: "Configuring a Remote Management Address for Outbound LLDP Advertisements")

Figure 15-12. Example of Per-Port Configuration Display (5300xl with Software Release E.10.x or Greater)

Configuring Global LLDP Packet Controls

The commands in this section configure the aspects of LLDP operation that apply the same to all ports in the switch.

Enabling or Disabling LLDP Operation on the Switch. Enabling LLDP operation (the default) causes the switch to:

- Use active, LLDP-enabled ports to transmit LLDP packets describing itself to neighbor devices.
- Add entries to its neighbors table based on data read from incoming LLDP advertisements.

Syntax [no] lldp run

*Enables or disables LLDP operation on the switch. The **no** form of the command, regardless of individual LLDP port configurations, prevents the switch from transmitting outbound LLDP advertisements, and causes the switch to drop all LLDP advertisements received from other devices. The switch preserves the current LLDP configuration when LLDP is disabled. After LLDP is disabled, the information in the LLDP neighbors database remains until it times-out. (Default: Enabled)*

For example, to disable LLDP on the switch:

```
ProCurve(config)# no lldp run
```

Changing the Packet Transmission Interval. This interval controls how often active ports retransmit advertisements to their neighbors.

Syntax lldp refresh-interval < 5 - 32768 >

Changes the interval between consecutive transmissions of LLDP advertisements on any given port. (Default: 30 seconds)

Note: The **refresh-interval** must be greater than or equal to ($4 \times \text{delay-interval}$). (The default **delay-interval** is 2). For example, with the default **delay-interval**, the lowest **refresh-interval** you can use is 8 seconds ($4 \times 2 = 8$). Thus, if you want a **refresh-interval** of 5 seconds, you must first change the delay interval to 1 (that is, $4 \times 1 < 5$). If you want to change the **delay-interval**, use the **setmib** command.

Changing the Time-to-Live for Transmitted Advertisements. The Time-to-Live value (in seconds) for all LLDP advertisements transmitted from a switch is controlled by the switch that generates the advertisement, and determines how long an LLDP neighbor retains the advertised data before

discarding it. The Time-to-Live value is the result of multiplying the **refresh-interval** by the **holdtime-multiplier** described below.

Syntax `lldp holdtime-multiplier < 2 - 10 >`

Changes the multiplier an LLDP switch uses to calculate the Time-to-Live for the LLDP advertisements it generates and transmits to LLDP neighbors. When the Time-to-Live for a given advertisement expires the advertised data is deleted from the neighbor switch's MIB. (Default: 4; Range: 2 - 10)

For example, if the refresh-interval on the switch is 15 seconds and the **holdtime-multiplier** is at the default, the Time-to-Live for advertisements transmitted from the switch is 60 seconds (4 x 15). To reduce the Time-to-Live, you could lower the **holdtime-interval** to 2, which would result in a Time-to-Live of 30 seconds.

```
ProCurve(config)# lldp holdtime-multiplier 2
```

Changing the Delay Interval Between Advertisements Generated by Value or Status Changes to the LLDP MIB. The switch uses a *delay-interval* setting to delay transmitting successive advertisements resulting from these LLDP MIB changes. If a switch is subject to frequent changes to its LLDP MIB, lengthening this interval can reduce the frequency of successive advertisements. The delay-interval can be changed using either an SNMP network management application or the CLI **setmib** command.

Syntax `setmib lldpTxDelay.0 -i < 1 - 8192 >`

Uses **setmib** to change the minimum time (delay-interval) any LLDP port will delay advertising successive LLDP advertisements due to a change in LLDP MIB content. (Default: 2; Range: 1 - 8192)

Note: The LLDP refresh-interval (transmit interval) must be greater than or equal to (4 x delay-interval). The switch does not allow increasing the delay interval to a value that conflicts with this relationship. That is, the switch displays **Inconsistent value** if (4 x delay-interval) exceeds the current transmit interval, and the command fails. Depending on the current refresh-interval setting, it may be necessary to increase the refresh-interval before using this command to increase the delay-interval.

For example, to change the delay-interval from 2 seconds to 8 seconds when the refresh-interval is at the default 30 seconds, you must first set the refresh-interval to a minimum of 32 seconds ($32 = 4 \times 8$).

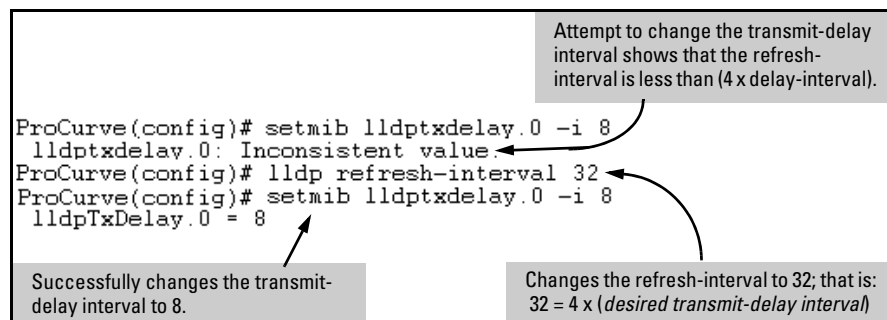


Figure 15-13. Example of Changing the Transmit-Delay Interval

Changing the Reinitialization Delay Interval. In the default configuration, a port receiving a **disable** command followed immediately by a **txonly**, **rxonly**, or **tx_rx** command delays reinitializing for two seconds, during which time LLDP operation remains disabled. If an active port is subjected to frequent toggling between the LLDP disabled and enabled states, LLDP advertisements are more frequently transmitted to the neighbor device. Also, the neighbor table in the adjacent device will change more frequently, as it deletes, then replaces LLDP data for the affected port which, in turn, generates SNMP traps (if trap receivers and SNMP notification are configured). All of this can unnecessarily increase network traffic. Extending the reinitialization-

delay interval delays the port's ability to reinitialize and generate LLDP traffic following an LLDP disable/enable cycle.

Syntax `setmib lldpReinitDelay.0 -i < 1 - 10 >`

*Uses **setmib** to change the minimum time (reinitialization delay interval) an LLDP port will wait before reinitializing after receiving an LLDP disable command followed closely by a `txonly` or `tx_rx` command. The delay interval commences with execution of the **lldp admin-status < port-list > disable** command. (Default: 2 seconds; Range: 1 - 10 seconds)*

For example, the following command changes the reinitialization delay interval to five seconds:

```
ProCurve(config)# setmib lldpreinitdelay.0 -i 5
```

Configuring SNMP Notification Support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices, and control the interval between successive notifications of data changes on the same neighbor.

Enabling LLDP Data Change Notification for SNMP Trap Receivers.

Syntax `[no] lldp enable-notification < port-list >`

Enables or disables each port in < port-list > for sending notification to configured SNMP trap receiver(s) if an LLDP data change is detected in an advertisement received on the port from an LLDP neighbor. (Default: Disabled)

For information on configuring trap receivers in the switch, refer to the chapter titled "Configuring for Network Management Applications" in the Management and Configuration Guide for your switch.

For example, this command enables SNMP notification on ports 1 - 5:

```
ProCurve(config)# lldp enable-notification 1-5
```

Changing the Minimum Interval for Successive Data Change Notifications for the Same Neighbor

LLDP trap notification is enabled on a port, a rapid succession of changes in LLDP information received in advertisements from one or more neighbors can generate a high number of traps. To reduce this effect, you can globally change the interval between successive notifications of neighbor data change.

Syntax `setmib lldpnotificationinterval.0 -i < 1 - 3600 >`

Globally changes the interval between successive traps generated by the switch. If multiple traps are generated in the specified interval, only the first trap will be sent. The remaining traps will be suppressed. (A network management application can periodically check the switch MIB to detect any missed change notification traps. Refer to IEEE P802.1AB or later for more information.) (Default: 5 seconds)

For example, the following command limits change notification traps from a particular switch to one per minute.

```
ProCurve(config)# setmib lldpnotificationinterval.0 -i 60  
lldpNotificationInterval.0 = 60
```

Configuring Per-Port Transmit and Receive Modes

These commands control advertisement traffic inbound and outbound on active ports.

Syntax `lldp admin-status < port-list > < txonly | rxonly | tx_rx | disable >`

With LLDP enabled on the switch in the default configuration, each port is configured to transmit and receive LLDP packets. These options enable you to control which ports participate in LLDP traffic and whether the participating ports allow LLDP traffic in only one direction or in both directions.

txonly: *Configures the specified port(s) to transmit LLDP packets, but block inbound LLDP packets from neighbor devices.*

rxonly: *Configures the specified port(s) to receive LLDP packets from neighbors, but block outbound packets to neighbors.*

tx_rx: *Configures the specified port(s) to both transmit and receive LLDP packets. (This is the default setting.)*

disable: *Disables LLDP packet transmit and receive on the specified port(s).*

Configuring Basic LLDP Per-Port Advertisement Content

In the default LLDP configuration, outbound advertisements from each port on the switch include both mandatory and optional data.

Mandatory Data. An active LLDP port on the switch always includes the mandatory data in its outbound advertisements. LLDP collects the mandatory data, and, except for the Remote Management Address, you cannot use LLDP commands to configure the actual data.

- Chassis Type (TLV subelement)
- Chassis ID (TLV)
- Port Type (TLV subelement)
- Port ID (TLV)
- Remote Management Address (TLV; actual IP address is a subelement that can be a default address or a configured address)

Configuring a Remote Management Address for Outbound LLDP Advertisements. This is an optional command you can use to include a specific IP address in the outbound LLDP advertisements for specific ports.

Syntax [no] lldp config < port-list > ipAddrEnable < ip-address >

*Replaces the default IP address for the port with an IP address you specify. This can be any IP address configured in a static VLAN on the switch, even if the port does not belong to the VLAN configured with the selected IP address. The **no** form of the command deletes the specified IP address. If there are no IP addresses configured as management addresses, then the IP address selection method returns to the default operation. (Default: The port advertises the IP address of the lowest-numbered VLAN (VID) to which it belongs. If there is no IP address configured on the VLAN(s) to which the port belongs, and the port is not configured to advertise an IP address from any other (static) VLAN on the switch, then the port advertises an address of 127.0.0.1.)*

Note: *This command does not accept either IP addresses acquired through DHCP or Bootp, or IP addresses that are not configured in a static VLAN on the switch*

For example, if port 3 belongs to a subnetted VLAN that includes an IP address of 10.10.10.100 and you wanted port 3 to use this secondary address in LLDP advertisements, you would need to execute the following command:

```
ProCurve(config)# lldp config 3 ipAddrEnable 10.10.10.100
```

Optional Data. You can configure an individual port or group of ports to exclude one or more of these data types from outbound LLDP advertisements. Note that optional data types, when enabled, are populated with data internal to the switch; that is, you cannot use LLDP commands to configure their actual content.

- port description (TLV)
- system name (TLV)
- system description (TLV)
- system capabilities (TLV)
 - system capabilities Supported (TLV subelement)
 - system capabilities Enabled (TLV subelement)
- port speed and duplex (TLV subelement)

Syntax: [no] lldp config < port-list > basicTlvEnable < TLV-Type >

port_descr

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the port.
(Default: Enabled)

system_name

For outbound LLDP advertisements, this TLV includes an alphanumeric string showing the system's assigned name.
(Default: Enabled)

system_descr

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the full name and version identification for the system's hardware type, software version, and networking application.
(Default: Enabled)

system_cap

*For outbound advertisements, this TLV includes a bitmask of supported system capabilities (device functions). Also includes information on whether the capabilities are enabled.
(Default: Enabled)*

For example, if you wanted to exclude the system name TLV from the outbound LLDP advertisements for all ports on a 3400cl-24G switch, you would use this command:

```
ProCurve(config)# no lldp config 1-24 basicTlvEnable  
system_name
```

If you later decided to reinstate the system name TLV on ports 1-5, you would use this command:

```
ProCurve(config)# lldp config 1-5 basicTlvEnable  
system_name
```

Configuring Support for Port Speed and Duplex Advertisements on the 5300xl and 4200vl Switches

This feature operates only on 5300xl switches running software release E.10.x or greater, and 4200vl switches.

This feature is enabled in the default LLDP-MED configuration on 5300xl switches running software release E.10.x or greater, and on 4200vl switches. It is optional for LLDP operation, but is *required* for LLDP-MED operation.

Port speed and duplex advertisements are supported on 5300xl switches and 4200vl switches to inform an LLDP endpoint and the switch port of each other's port speed and duplex configuration and capabilities. Configuration mismatches between a switch port and an LLDP endpoint can result in excessive collisions and voice quality degradation. LLDP enables discovery of such mismatches by supporting SNMP access to the switch MIB for comparing the current switch port and endpoint settings. (Changing a current device configuration to eliminate a mismatch requires intervention by the system operator.)

Syntax: [no] lldp config < port-list > dot3TlvEnable macphy_config

Note: *This command applies only to 5300xl switches running software release E.10.xx or greater, and 4200vl switches.*

For outbound advertisements, this TLV includes the (local) switch port's current speed and duplex settings, the range of speed and duplex settings the port supports, and the method required for reconfiguring the speed and duplex settings on the device (auto-negotiation during link initialization, or manual configuration).

*Using SNMP to compare local and remote information can help in locating configuration mismatches.
(Default: Enabled)*

Note: *For LLDP operation, this TLV is optional. For LLDP-MED operation, this TLV is mandatory.*

As mentioned above, an SNMP network management application can be used to compare the port speed and duplex data configured in the switch and advertised by the LLDP endpoint. You can also use the CLI to display this information. For more on using the CLI to display port speed and duplex information, refer to “Displaying the Current Port Speed and Duplex Configuration on a Switch Port” on page 15-64.

LLDP-MED (Media-Endpoint-Discovery) for the 5300xl and 4200vl Switches

As of October 2006, LLDP-MED operates only on 5300xl switches running software release E.10.x or greater, and 4200vl switches.

LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED uses the standard LLDP commands described earlier in this section, with some extensions, and also introduces new commands unique to LLDP-MED operation. The **show** commands described elsewhere in this section are applicable to both LLDP and LLDP-MED operation. LLDP-MED benefits include:

- plug-and-play provisioning for MED-capable, VoIP endpoint devices
- simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network
- automatic deployment of convergence network policies (voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority)
- configurable endpoint location data to support the Emergency Call Service (ECS) (such as Enhanced 911 service, 999, 112)
- detailed VoIP endpoint data inventory readable via SNMP from the switch
- Power over Ethernet (PoE) status and troubleshooting support via SNMP
- support for IP telephony network troubleshooting of call quality issues via SNMP

This section describes how to configure and use LLDP-MED features in the 5300xl switches running software release E.10.x or greater, and 4200vl switches, to support VoIP network edge devices (Media Endpoint Devices) such as:

- IP phones
- voice/media gateways
- media servers

- IP communications controllers
- other VoIP devices or servers

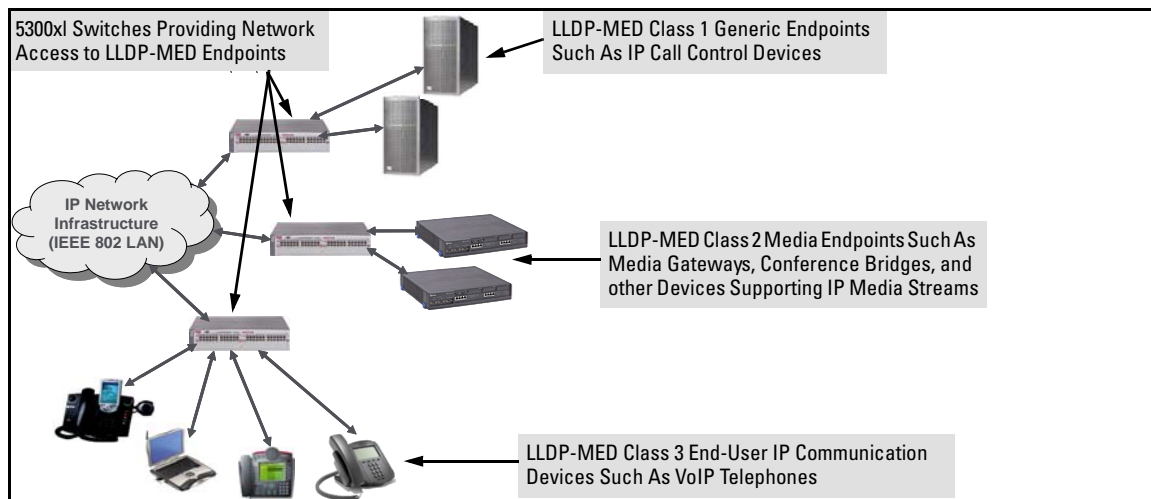


Figure 15-14. Example of LLDP-MED Network Elements

LLDP-MED Endpoint Support. LLDP-MED interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- able to autonegotiate speed and duplex configuration with the switch
- able to use the following network policy elements configured on the client port
 - voice VLAN ID
 - 802.1p (Layer 2) QoS
 - Diffserv codepoint (DSCP) (Layer 3) QoS
- discover and advertise device location data learned from the switch
- support emergency call service (ECS—such as E911, 999, and 112)
- advertise device information for the device data inventory collected by the switch, including:
 - hardware revision
 - serial number
 - asset ID
 - firmware revision
 - manufacturer name
 - software revision
 - model name

- provide information on network connectivity capabilities (for example, a multi-port VoIP phone with Layer 2 switch capability)
- support the fast start capability

Note

LLDP-MED is intended for use with VoIP endpoints, and is not designed to support links between network infrastructure devices, such as switch-to-switch or switch-to-router links.

LLDP-MED Endpoint Device Classes. LLDP-MED endpoint devices are, by definition, located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (Generic Endpoint Devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (Media Endpoint Devices): These devices offer all Class 1 features plus media streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.
- Class 3 (Communication Devices): These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

LLDP-MED Operational Support on the 5300xl and 4200v1 Switches.

The 4200v1 switches, and the 5300xl switches beginning with software release E.10.xx, offer two configurable TLVs supporting MED-specific capabilities:

- medTlvEnable (for per-port enabling or disabling of LLDP-MED operation)
- medPortLocation (for configuring per-port location or emergency call data)

Note

LLDP-MED operation also requires the port speed and duplex TLV (dot3TlvEnable; page 15-46), which is enabled in the default configuration.

LLDP-MED Topology Change Notification

This optional feature provides information an SNMP application can use to track LLDP-MED connects and disconnects.

Syntax: `lldp top-change-notify < port-list >`

Topology change notification, when enabled on an LLDP port, causes the switch to send an SNMP trap if it detects LLDP-MED endpoint connection or disconnection activity on the port, or an age-out of the LLDP-MED neighbor on the port. The trap includes the following information:

- *the port number (internal) on which the activity was detected (For more in internal port numbers, refer to “Determining the 5300xl Port Number Included in Topology Change Notification Traps” on page 15-71.)*
- *the LLDP-MED class of the device detected on the port (“LLDP-MED Endpoint Device Classes” on page 15-49.)*

The **show running** command shows whether the topology change notification feature is enabled or disabled. For example, if ports A1-A10 have topology change notification enabled, the following entry appears in the **show running** output:

```
lldp top-change-notify A1-A10
```

(Default: Disabled)

Note: To send traps, this feature requires access to at least one SNMP server. For information on configuring traps, go to the chapter titled “Configuring for Network Management Applications” in the Management and Configuration Guide for your switch, and refer to one of the following sections:

- SNMPv1 and SNMPv2c Trap Features
- SNMPv3 Notification and Traps

Also, if a detected LLDP-MED neighbor begins sending advertisements without LLDP-MED TLVs, the switch sends a top-change-notify trap.

Note

Topology change notifications provide one method for monitoring system activity. However, because SNMP normally employs UDP, which does not guarantee datagram delivery, topology change notification should not be relied upon as the sole method for monitoring critical endpoint device connectivity.

LLDP-MED Fast Start Control

Syntax: `lldp fast-start-count < 1 - 10 >`

*An LLDP-MED device connecting to a switch port may use the data contained in the MED TLVs from the switch to configure itself. However, the **lldp refresh-interval** setting (default: 30 seconds) for transmitting advertisements can cause an unacceptable delay in MED device configuration. To support rapid LLDP-MED device configuration, the **lldp fast-start-count** command temporarily overrides the **refresh-interval** setting for the **fast-start-count** advertisement interval. This results in the port initially advertising LLDP-MED at a faster rate for a limited time. Thus, when the switch detects a new LLDP-MED device on a port, it transmits one LLDP-MED advertisement per second out the port for the duration of the **fast-start-count** interval. In most cases, the default setting should provide an adequate **fast-start-count** interval.*

(Range: 1 - 10 seconds; Default: 5 seconds)

Note: This global command applies only to ports on which a new LLDP-MED device is detected. It does not override the refresh-interval setting on ports where non-MED devices are detected.

Advertising Device Capability, Network Policy, PoE Status and Location Data

The `medTlvEnable` option on the switch is enabled in the default configuration and supports the following LLDP-MED TLVs:

- LLDP-MED capabilities: This TLV enables the switch to determine:
 - whether a connected endpoint device supports LLDP-MED
 - which specific LLDP-MED TLVs the endpoint supports
 - the device class (1, 2, or 3) for the connected endpoint

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

- network policy operating on the port to which the endpoint is connected (VLAN, Layer 2 QoS, Layer 3 QoS)
- PoE (MED Power-over-Ethernet)
- physical location data — page 56

Note

LLDP-MED operation requires the `macphy_config` TLV subelement—enabled by default—that is optional for IEEE 802.1AB LLDP operation. Refer to the **`dot3TlvEnable macphy_config`** command on page 15-46.

Network Policy Advertisements. Network policy advertisements are intended for real-time voice and video applications, and include these TLV subelements:

- Layer 2 (802.1p) QoS
- Layer 3 DSCP (diffserv code point) QoS
- Voice VLAN ID (VID)

VLAN Operating Rules. These rules affect advertisements of VLANs in network policy TLVs:

- The VLAN ID TLV subelement applies only to a VLAN configured for voice operation (**`vlan < vid > voice`**).
- If there are multiple voice VLANs configured on a port, LLDP-MED advertises the voice VLAN having the lowest VID.

- The voice VLAN port membership configured on the switch can be tagged or untagged. However, if the LLDP-MED endpoint expects a tagged membership when the switch port is configured for untagged, or the reverse, then a configuration mismatch results. (Typically, the endpoint expects the switch port to have a tagged voice VLAN membership.)
- If a given port does not belong to a voice VLAN, then the switch does not advertise the VLAN ID TLV through this port.

Policy Elements. These policy elements may be statically configured on the switch or dynamically imposed during an authenticated session using a RADIUS server and 802.1X or MAC authentication. (Web authentication does not apply to VoIP telephones and other telecommunications devices that are not capable of accessing the switch through a Web browser.) The QoS and voice VLAN policy elements can be statically configured with the following CLI commands:

```
vlan < vid > voice  
vlan < vid > < tagged | untagged > < port-list >  
int < port-list > qos priority < 0 - 7 >  
vlan < vid > qos dscp < codepoint >
```

Notes

A codepoint must have an 802.1p priority before you can configure it for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows **No Override** in the **Priority** column of the DSCP policy table (display with **show qos-dscp map**, then use **qos-dscp map < codepoint > priority < 0 - 7 >** to configure a priority before proceeding. For more on this topic, refer to the chapter titled “Quality of Service (QoS): Managing Bandwidth More Effectively” in the *Advanced Traffic Management Guide* for your switch.

Enabling or Disabling medTlvEnable on 5300xl or 4200vl Switches.

In the default LLDP-MED configuration, the TLVs controlled by medTlvEnable are enabled.

Syntax: [no] lldp config < port-list > medTlvEnable < medTlv >

- *Enables or disables advertisement of the following TLVs on the specified ports:*
 - *device capability TLV*
 - *configured network policy TLV*
 - *configured location data TLV (Refer to “Configuring Location Data for LLDP-MED Devices” on page 15-56.)*
 - *current PoE status TLV*
- (Default: All of the above TLVs are enabled.)
- *Helps to locate configuration mismatches by allowing use of an SNMP application to compare the LLDP-MED configuration on a port with the LLDP-MED TLVs advertised by a neighbor connected to that port.*

capabilities

This TLV enables the switch to determine:

- *which LLDP-MED TLVs a connected endpoint can discover*
- *the device class (1, 2, or 3) for the connected endpoint*

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

(Default: enabled)

Note: *This TLV cannot be disabled unless the network_policy, poe, and location_id TLVs are already disabled.*

network-policy

This TLV enables the switch port to advertise its configured network policies (voice VLAN, Layer 2 QoS, Layer 3 QoS), and allows LLDP-MED endpoint devices to auto-configure the voice network policy advertised by the switch. This also enables the use of SNMP applications to troubleshoot statically configured endpoint network policy mismatches.

(Default: Enabled)

Notes: Network policy is only advertised for ports that are configured as members of the voice VLAN. If the port belongs to more than one voice VLAN, then the voice VLAN with the lowest-numbered VID is selected as the VLAN for voice traffic. Also, this TLV cannot be enabled unless the capability TLV is already enabled.

For more information, refer to “Network Policy Advertisements” on page 15-52

location_id

This TLV enables the switch port to advertise its configured location data (if any). For more on configuring location data, refer to “Configuring Location Data for LLDP-MED Devices”.

(Default: Enabled)

Note: When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.

poe

This TLV enables the switch port to advertise its current PoE (Power over Ethernet) state and to read the PoE requirements advertised by the LLDP-MED endpoint device connected to the port.

(Default: Enabled)

Note: When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.

For more on this topic, refer to “PoE Advertisements”, below.

PoE Advertisements. These advertisements inform an LLDP-MED endpoint of the power (PoE) configuration on switch ports. Similar advertisements from an LLDP-MED endpoint inform the switch of the endpoint's power needs and provide information that can be used to identify power priority mismatches.

Power-over-Ethernet TLVs include the following power data:

- **power type:** indicates whether the device is a power-sourcing entity (PSE) or a powered device (PD). Ports on the J8161A PoE xl module are PSE devices. A MED-capable VoIP telephone is a PD.
- **power source:** indicates the source of power in use by the device. Power sources for powered devices (PDs) include PSE, local (internal), and PSE/local. The 5300xl switches advertise Unknown.
- **power priority:** indicates the power priority configured on the switch (PSE) port or the power priority configured on the MED-capable endpoint.
- **power value:** indicates the total power in watts that a switch port (PSE) can deliver at a particular time, or the total power in watts that the MED endpoint (PD) requires to operate.

To display the current power data for an LLDP-MED device connected to a port, use the following command:

show lldp info remote-device < port-list >

For more on this command, refer to page 66.

To display the current PoE configuration on the switch, use the following commands:

show power brief < port-list >

show power < port-list >

For more on PoE configuration and operation, refer to the chapter titled “Power Over Ethernet (PoE) Operation for the Series 5300xl Switches” in the *Management and Configuration Guide* for your switch.

Configuring Location Data for LLDP-MED Devices

You can configure a switch port to advertise location data for the switch itself, the physical wall-jack location of the endpoint (recommended), or the location of a DHCP server supporting the switch and/or endpoint. You also have the option of configuring these different address types:

- **civic address:** physical address data such as city, street number, and building information

- **ELIN (Emergency Location Identification Number):** an emergency number typically assigned to MLTS (Multiline Telephone System Operators) in North America
- **coordinate-based location:** attitude, longitude, and altitude information (Requires configuration via an SNMP application.)

Syntax: [no] lldp config < port-list > medPortLocation < Address-Type >

*Configures location or emergency call data the switch advertises per port in the **location_id** TLV. This TLV is for use by LLDP-MED endpoints employing location-based applications.*

Note: *The switch allows one medPortLocation entry per port (without regard to type). Configuring a new medPortLocation entry of any type on a port replaces any previously configured entry on that port.*

civic-addr < COUNTRY-STR > < WHAT > < CA-TYPE > < CA-VALUE > ...
[< CA-TYPE > < CA-VALUE >] ... [< CA-TYPE > < CA-VALUE >]

This command enables configuration of a physical address on a switch port, and allows up to 75 characters of address information.

COUNTRY-STR: *A two-character country code, as defined by ISO 3166. Some examples include **FR** (France), **DE** (Germany), and **IN** (India). This field is required in a **civic-addr** command. (For a complete list of country codes, visit www.iso.org on the world wide web.)*

WHAT: *A single-digit number specifying the type of device to which the location data applies:*

0: *Location of DHCP server*

1: *Location of switch*

2: *Location of LLDP-MED endpoint (recommended application)*

*This field is required in a **civic-addr** command.*

Type/Value Pairs (CA-TYPE and CA-VALUE): *This is a series of data pairs, each composed of a location data “type” specifier and the corresponding location data for that type. That is, the first value in a pair is expected to be the civic address “type” number (**CA-TYPE**), and the second value in a pair is expected to be the corresponding civic address data (**CA-VALUE**). For example, if the **CA-TYPE** for “city name” is “3”, then the type/value pair to define the city of Paris is “3 Paris”. Multiple type/value pairs can be entered in any order, although it is recommended that multiple pairs be entered in ascending order of the **CA-TYPE**.*

*When an emergency call is placed from a properly configured class 3 endpoint device to an appropriate PSAP, the country code, device type, and type/value pairs configured on the switch port are included in the transmission. The “type” specifiers are used by the PSAP to identify and organize the location data components in an understandable format for response personnel to interpret. A **civic-addr** command requires a minimum of one type/value pair, but typically includes multiple type/value pairs as needed to configure a complete set of data describing a given location.*

CA-TYPE: *This is the first entry in a type/value pair, and is a number defining the type of data contained in the second entry in the type/value pair (**CA-VALUE**). Some examples of **CA-TYPE** specifiers include:*

- 3 = city
- 6 = street (name)
- 25 = building name

(Range: 0 - 255)

*For a sample listing of **CA-TYPE** specifiers, refer to table 15-5 on page 15-60.*

CA-VALUE: *This is the second entry in a type/value pair, and is an alphanumeric string containing the location information corresponding to the immediately preceding **CA-TYPE** entry. Strings are delimited by either blank spaces, single quotes ('...'), or double quotes (“...”). Each string should represent a specific data type in a set of unique type/value pairs comprising the description of a location, and each string must be preceded by a **CA-TYPE** number identifying the type of data in the string.*

Note: *A 5300xl port allows one instance of any given **CA-TYPE**. For example, if a type/value pair of **6 Atlantic** (to specify “Atlantic” as a street name) is configured on port A5 and later another type/value pair of **6 Pacific** is configured on the same port, then **Pacific replaces Atlantic** in the civic address location configured for port A5.*

elin-addr < emergency-number >

*This feature is intended for use in Emergency Call Service (ECS) applications to support class 3 LLDP-MED VoIP telephones connected to a 5300xl switch in a multiline telephone system (MLTS) infrastructure. An ELIN (Emergency Location Identification Number) is a valid North American Numbering Plan (NANP) format telephone number assigned to MLTS operators in North America by the appropriate authority. The ELIN is used to route emergency (E911) calls to a Public Safety Answering Point (PSAP).
(Range: 1-15 numeric characters)*

Configuring Coordinate-Based Locations. Latitude, longitude, and altitude data can be configured per switch port using an SNMP management application. For more information, refer to the documentation provided with the application. A further source of information on this topic is *RFC 3825-Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information*.

Note

Endpoint use of data from a medPortLocation TLV sent by the switch is device-dependent. Refer to the documentation provided with the endpoint device.

Table 15-5. Some Location Codes Used in CA-TYPE Fields*

Location Element	Code	Location Element	Code
national subdivision	1	street number	19
regional subdivision	2	additional location data	22
city or township	3	unit or apartment	26
city subdivision	4	floor	27
street	6	room number	28
street suffix	18		
*The code assignments in this table are examples from a work-in-progress (the internet draft titled “Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information draft-ietf-geopriv-dhcp-civil-06” dated May 30, 2005.) For the actual codes to use, contact the PSAP or other authority responsible for specifying the civic addressing data standard for your network.			

Example of a Location Configuration on a 5300xl Switch Port.

Suppose a system operator wanted to configure the following information as the civic address for a telephone connected to her company’s network through port A2 of a 5300xl switch at the following location:

Description	CA-Type	CA-VALUE
national subdivision	1	CA
city	3	Widgitville
street	6	Main
street number	19	1433
unit	26	Suite 4-N
floor	27	4
room number	28	N4-3

Figure 15-15 shows the commands for configuring and displaying the above data.

```
ProCurve(config)# lldp config a2 medportlocation civic-addr US 2 1 CA 3 Widgeville 6 Main 19 1433 26 Suite_4-N 27 4 28 N4-3
ProCurve(config)# show lldp config a2

LLDP Port Configuration Detail

Port : A2
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

Country Name      : US
What              : 2
Ca-Type           : 1
Ca-Length         : 2
Ca-Value          : CA
Ca-Type           : 3
Ca-Length         : 11
Ca-Value          : Widgeville
Ca-Type           : 6
Ca-Length         : 4
Ca-Value          : Main
Ca-Type           : 19
Ca-Length         : 4
Ca-Value          : 1433
Ca-Type           : 26
Ca-Length         : 9
Ca-Value          : Suite_4-N
Ca-Type           : 27
Ca-Length         : 1
Ca-Value          : 4
Ca-Type           : 28
Ca-Length         : 4
Ca-Value          : N4-3
```

Figure 15-15. Example of a Civic Address Configuration

Displaying Advertisement Data

Command	Page
show lldp info local-device	below
walkmib lldpXdot3LocPortOperMauType	
show lldp info remote-device	15-65
walkmib lldpXdot3RemPortAutoNegAdvertisedCap	
show lldp info stats	15-68

Displaying Switch Information Available for Outbound Advertisements

These commands display the current switch information that will be used to populate outbound LLDP advertisements.

Syntax `show lldp info local-device [port-list]`

Without the [port-list] option, this command displays the global switch information and the per-port information currently available for populating outbound LLDP advertisements.

With the [port-list] option, this command displays only the following port-specific information that is currently available for outbound LLDP advertisements on the specified ports:

- **PortType**
- **PortId**
- **PortDesc**

Note: *This command displays the information available on the switch. Use the `lldp config < port-list >` command to change the selection of information that is included in actual outbound advertisements. In the default LLDP configuration, all information displayed by this command is transmitted in outbound advertisements.*

For example, in the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in figure 15-16 on page 15-64.

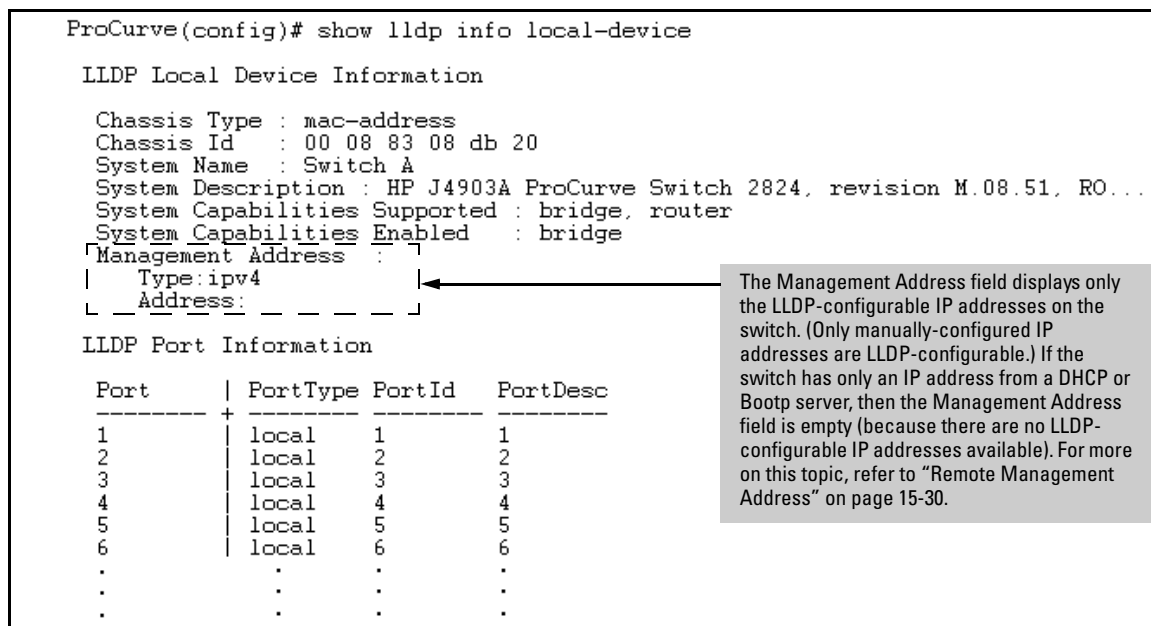


Figure 15-16. Example of Displaying the Global and Per-Port Information Available for Outbound Advertisements

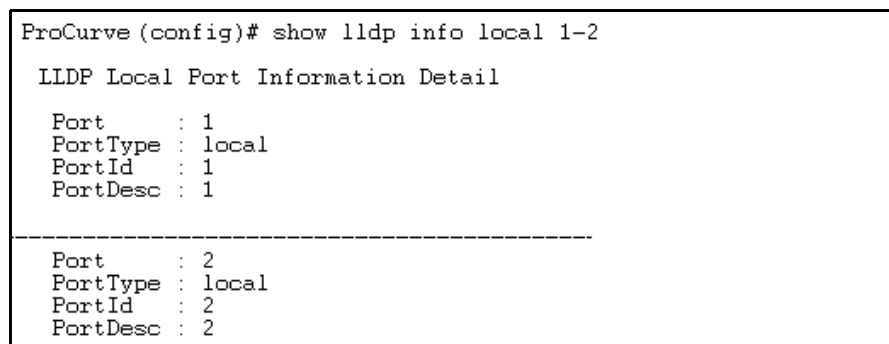


Figure 15-17. Example of the Default Per-Port Information Content for Ports 1 and 2

Displaying the Current Port Speed and Duplex Configuration on a Switch Port. Port speed and duplex information for a switch port and a connected LLDP-MED endpoint can be compared for configuration mismatches by using an SNMP application. You can also use the switch CLI to display this information, if necessary. The following two commands provide methods for displaying speed and duplex information for switch ports. For

information on displaying the currently configured port speed and duplex on an LLDP-MED endpoint, refer to “Displaying the Current Port Speed and Duplex Configuration on a Switch Port” on page 15-64.

Syntax: show interfaces brief < port-list >

*Includes port speed and duplex configuration in the **Mode** column of the resulting display.*

(This command is available on all switch models covered by this guide.)

Displaying Advertisements Currently in the Neighbors MIB. These commands display the content of the inbound LLDP advertisements received from other LLDP devices.

Syntax show lldp info remote-device [*port-list*]

*Without the [*port-list*] option, this command provides a global list of the individual devices it has detected by reading LLDP advertisements. Discovered devices are listed by the inbound port on which they were discovered. Multiple devices listed for a single port indicates one of the following:*

- **5300xl and 4200xl Switches:** Multiple devices listed for the same port indicates that such devices are connected to the switch through a hub.
- **3400cl/6400cl Switches:** Multiple devices listed for the same port indicate that multiple devices are connected to the switch through a hub.

Discovering the same device on multiple ports indicates that the remote device may be connected to the switch in one of the following ways:

- *Through different VLANs using separate links. (This applies to switches that use the same MAC address for all configured VLANs.)*
- *Through different links in the same trunk.*
- *Through different links using the same VLAN. (In this case, spanning-tree should be invoked to prevent a network topology loop. Note that LLDP packets travel on links that spanning-tree blocks for other traffic types.)*

*With the [*port-list*] option, this command provides a listing of the LLDP data that the switch has detected in advertisements received on the specified ports.*

For descriptions of the various types of information displayed by these commands, refer to Table 15-3 on page 15-29.

```
ProCurve Switch 3400cl-24G# show lldp info remote
```

LLDP Remote Devices Information

LocalPort	ChassisId	PortId	PortName	SysName
1	00 11 85 c6 54 60	17	17	ProCurve Switch ...
2	00 11 85 cf 66 80	33	33	ProCurve Switch ...

Note: In software releases earlier than M_08_06x (for the 3400cl switches only), a **Port Type** column appears with this command instead of the **PortId**, **PortName** columns shown in this figure.

Note: In software release E.10.x and greater for the 5300xl switches, and for 4200vl switches, the **PortName** column heading appears as **PortDescr**.

Figure 15-18. Example of a Global Listing of Discovered Devices

```
ProCurve(config)# show lldp info remote-device a2
```

LLDP Remote Device Information Detail

```
Local Port      : A2
ChassisType     : network-address
ChassisId       : 0f ff 7a 5c
PortType        : mac-address
PortId          : 08 00 0f 14 de f2
SysName         : regDN 3004,<IP-Phone-Data>
System Descr    : regDN 3004,<IP-Phone-Data>,h/w rev 0,ASIC rev 0,f/w Boot FW...
PortDescr       : LAN port
```

```
System Capabilities Supported : bridge, telephone
System Capabilities Enabled   : bridge, telephone
```

Remote Management Address

MED Information Detail

```
EndpointClass   :Class3
Media Policy Vlan id :10
Media Policy Priority :7
Media Policy Dscp   :44
Media Policy Tagged  :False
Poe Device Type    :PD
Power Requested     :47
Power Source        :Unknown
Power Priority       :High
```

Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.

Figure 15-19. Example of a 5300xl LLDP-MED Listing of an Advertisement Received From an LLDP-MED (VoIP Telephone) Source

Displaying LLDP Statistics

LLDP statistics are available on both a global and a per-port levels. Rebooting the switch resets the LLDP statistics counters to zero. Disabling the transmit and/or receive capability on a port “freezes” the related port counters at their current values.

Syntax `show lldp stats [port-list]`

The global LLDP statistics command displays an overview of neighbor detection activity on the switch, plus data on the number of frames sent, received, and discarded per-port. The per-port LLDP statistics command enhances the list of per-port statistics provided by the global statistics command with some additional per-port LLDP statistics.

Global LLDP Counters:

Neighbor Entries List Last Updated: *Shows the elapsed time since a neighbor was last added or deleted.*

New Neighbor Entries Count: *Shows the total of new LLDP neighbors detected since the last switch reboot. Disconnecting, then reconnecting a neighbor increments this counter.*

Neighbor Entries Deleted Count: *Shows the number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports. For example, if the admin status for port on a neighbor device changes from **tx_rx** or **txonly** to **disabled** or **rxonly**, then the neighbor device sends a “shutdown” packet out the port and ceases transmitting LLDP frames out that port. The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter.*

Neighbor Entries Dropped Count: *Shows the number of valid LLDP neighbors the switch detected, but could not add. This can occur, for example, when a new neighbor is detected when the switch is already supporting the maximum number of neighbors. Refer to “Neighbor Maximum” on page 15-70.*

Neighbor Entries AgeOut Count: *Shows the number of LLDP neighbors dropped on all ports due to Time-to-Live expiring.*

— Continued on the next page. —

— Continued from the preceding page. —

Per-Port LLDP Counters:

NumFramesRecvd: *Shows the total number of valid, inbound LLDP advertisements received from any neighbor(s) on <port-list>. Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources.*

NumFramesSent: *Shows the total number of LLDP advertisements sent from <port-list>.*

NumFramesDiscarded: *Shows the total number of inbound LLDP advertisements discarded by <port-list>. This can occur, for example, when a new neighbor is detected on the port, but the switch is already supporting the maximum number of neighbors. Refer to “Neighbor Maximum” on page 15-70. This can also be an indication of advertisement formatting problems in the neighbor device.*

Frames Invalid: *Shows the total number of invalid LLDP advertisements received on the port. An invalid advertisement can be caused by header formatting problems in the neighbor device.*

TLVs Unrecognized: *Shows the total number of LLDP TLVs received on a port with a type value in the reserved range. This could be caused by a basic management TLV from a later LLDP version than the one currently running on the switch.*

TLVs Discarded: *Shows the total number of LLDP TLVs discarded for any reason. In this case, the advertisement carrying the TLV may be accepted, but the individual TLV was not usable.*

Neighbor Ageouts: *Shows the number of LLDP neighbors dropped on the port due to Time-to-Live expiring.*

```
ProCurve(config)# show lldp stats
```

LLDP Device Statistics

Neighbor Entries List Last Updated : 2 hours
New Neighbor Entries Count : 20
Neighbor Entries Deleted Count : 20
Neighbor Entries Dropped Count : 0
Neighbor Entries AgeOut Count : 20

LLDP Port Statistics

Port	NumFramesRecvd	NumFramesSent	NumFramesDiscarded
1	628	316	0
2	21	12	0
3	0	252	0
4	446	226	0
5	0	0	0
6	0	0	0
.	.	.	.
.	.	.	.
.	.	.	.

Counters showing frames sent on a port but no frames received on that port indicates an active link with a device that either has LLDP disabled on the link or is not LLDP-aware.

Figure 15-20. Example of a Global LLDP Statistics Display

```
ProCurve(config)# show lldp stats 1
```

LLDP Port Statistics Detail

PortName : 1
Frames Discarded : 0
Frames Invalid : 0
Frames Received : 658
Frames Sent : 331
TLVs Unrecognized : 0
TLVs Discarded : 0
Neighbor Ageouts : 0

Figure 15-21. Example of a Per-Port LLDP Statistics Display

LLDP Operating Notes

Neighbor Maximum. The neighbors table in the switch supports as many neighbors as there are ports on the switch. The switch can support multiple neighbors connected through a hub on a given port, but if the switch neighbor maximum is reached, advertisements from additional neighbors on the same or other ports will not be stored in the neighbors table unless some existing neighbors time-out or are removed.

LLDP Packet Forwarding: An 802.1D-compliant switch does not forward LLDP packets, regardless of whether LLDP is globally enabled or disabled on the switch.

One IP Address Advertisement Per-Port: LLDP advertises only one IP address per-port, even if multiple IP addresses are configured by `lldp config <port-list> ipAddrEnable` on a given port.

802.1Q VLAN Information. LLDP packets do not include 802.1Q header information, and are always handled as untagged packets.

Effect of 802.1X Operation. If 802.1X port security is enabled on a port and a connected device is not authorized, LLDP packets are not transmitted or received on that port. Any neighbor data stored in the neighbor MIB for that port prior to the unauthorized device connection remains in the MIB until it ages out. If an unauthorized device later becomes authorized, LLDP transmit and receive operation resumes.

Neighbor Data Can Remain in the Neighbor Database After the Neighbor Is Disconnected. After disconnecting a neighbor LLDP device from the switch, the neighbor can continue to appear in the switch's neighbor database for an extended period if the neighbor's **holdtime-multiplier** is high; especially if the **refresh-interval** is large. Refer to "Changing the Time-to-Live for Transmitted Advertisements" on page 15-38.

Mandatory TLVs. All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

Determining the 5300xl Port Number Included in Topology Change Notification Traps. Enabling topology change notification on a 5300xl switch port and then connecting or disconnecting an LLDP-MED endpoint on that port causes the switch to send an SNMP trap to notify the designated management station(s). The port number included in the trap corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number appearing in an SNMP trap, use the **walkmib ifDescr** command, as shown in the following figure:

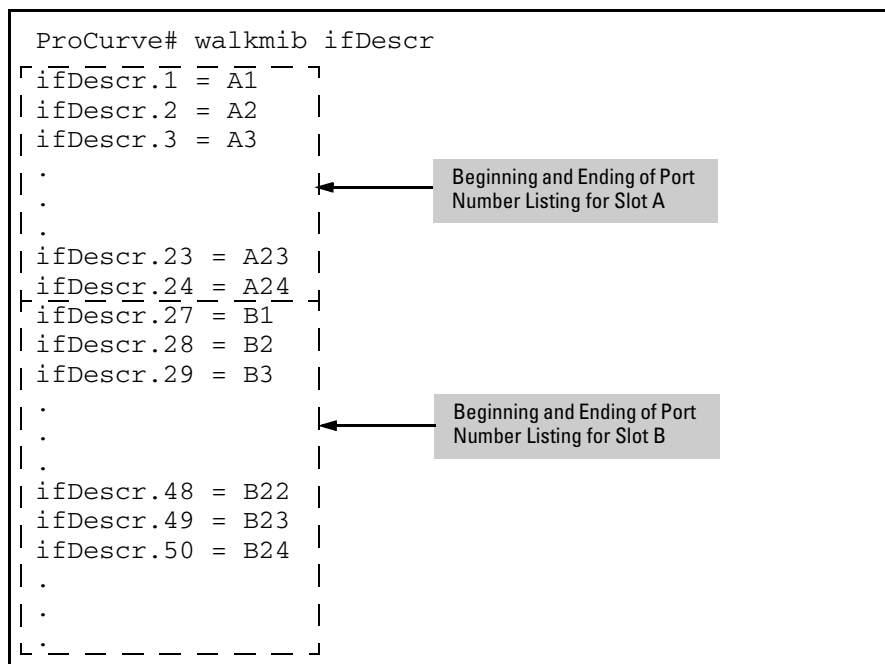


Figure 15-22. Matching Internal Port Numbers to External Slot/Port Numbers

LLDP and CDP Data Management

This section describes points to note regarding LLDP (Link-Layer Discovery Protocol) and CDP (Cisco Discovery Protocol) data received by the switch from other devices. LLDP operation includes both transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices. CDP operation is limited to reading incoming CDP packets from neighbor devices. (ProCurve switches do not generate CDP packets.)

LLDP and CDP Neighbor Data

With both LLDP and (read-only) CDP enabled on a switch port, the port can read both LLDP and CDP advertisements, and stores the data from both types of advertisements in its neighbor database. (The switch only *stores* CDP data that has a corresponding field in the LLDP neighbor database.) The neighbor database itself can be read by either LLDP or CDP methods or by using the **show lldp** commands. Take note of the following rules and conditions:

- If the switch receives both LLDP and CDP advertisements on the same port from the same neighbor the switch stores this information as two separate entries if the advertisements have differences chassis ID and port ID information.
- If the chassis and port ID information are the same, the switch stores this information as a single entry. That is, LLDP data overwrites the corresponding CDP data in the neighbor database if the chassis and port ID information in the LLDP and CDP advertisements received from the same device is the same.
- Data read from a CDP packet does not support some LLDP fields, such as “System Descr”, “SystemCapSupported”, and “ChassisType”. For such fields, LLDP assigns relevant default values. Also:
 - The LLDP “System Descr” field maps to CDP’s “Version” and “Platform” fields.
 - The switch assigns “ChassisType” and “PortType” fields as “local” for both the LLDP and the CDP advertisements it receives.
 - Both LLDP and CDP support the “System Capability” TLV. However, LLDP differentiates between what a device is capable of supporting and what it is actually supporting, and separates the two types of information into subelements of the System Capability TLV. CDP has only a single field for this data. Thus, when CDP System Capability data is mapped to LLDP, the same value appears in both LLDP System Capability fields.
 - System Name and Port Descr are not communicated by CDP, and thus are not included in the switch’s Neighbors database.

Note

Because ProCurve switches do not generate CDP packets, they are not represented in the CDP data collected by any neighbor devices running CDP.

A switch with CDP disabled forwards the CDP packets it receives from other devices, but does not store the CDP information from these packets in its own MIB.

LLDP data transmission/collection and CDP data collection are both enabled in the switch’s default configuration. In this state, an SNMP network management application designed to discover devices running either CDP or LLDP can retrieve neighbor information from the switch regardless of whether LLDP or CDP is used to collect the device-specific information.

Protocol State	Packet Generation	Inbound Data Management	Inbound Packet Forwarding
CDP Enabled ¹	n/a	Store inbound CDP data.	No forwarding of inbound CDP packets.
CDP Disabled	n/a	No storage of CDP data from neighbor devices.	Floods inbound CDP packets from connected devices to outbound ports.
LLDP Enabled ¹	Generates and transmits LLDP packets out all ports on the switch.	Store inbound LLDP data.	No forwarding of inbound LLDP packets.
LLDP Disabled	No packet generation.	No storage of LLDP data from neighbor devices.	No forwarding of inbound LLDP packets.

¹Both CDP data collection and LLDP transmit/receive are enabled in the default configuration. If a switch receives CDP packets and LLDP packets from the same neighbor device on the same port, it stores and displays the two types of information separately if the chassis and port ID information in the two types of advertisements is different. In this case, if you want to use only one type of data from a neighbor sending both types, disable the unwanted protocol on either the neighbor device or on the switch. However, if the chassis and port ID information in the two types of advertisements is the same, the LLDP information overwrites the CDP data for the same neighbor device on the same port.

CDP Operation and Commands

By default the switches covered by this guide have CDP enabled on each port. This is a read-only capability, meaning that the switch can receive and store information about adjacent CDP devices but does not generate CDP packets.

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received (and does not forward the packet). The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet, and cannot be controlled in the switch receiving the packet.) A switch reviews the list of CDP neighbor entries every three seconds, and purges any expired entries.

Command	Page
show cdp	15-75
show cdp neighbors [< port-list > detail] [detail < port-list >]	15-76
[no] cdp run	15-77
[no] cdp enable < port-list >	15-77

Note For details on how to use an SNMP utility to retrieve information from the switch's CDP Neighbors table maintained in the switch's MIB (Management Information Base), refer to the documentation provided with the particular SNMP utility.

Viewing the Switch's Current CDP Configuration. CDP is shown as enabled/disabled both globally on the switch and on a per-port basis.

Syntax: show cdp
Lists the switch's global and per-port CDP configuration.

The following example shows the default CDP configuration.

```
ProCurve(config)# show cdp
Global CDP information
  Enable CDP [Yes] : Yes

Port  CDP
----  -
A1    enabled
A2    enabled
A3    enabled
.      .
.      .
.      .
```

Figure 15-23. Example of Show CDP with the Default CDP Configuration

Viewing the Switch’s Current CDP Neighbors Table. Devices are listed by the port on which they were detected.

Syntax: show cdp neighbors

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device’s CDP packet.

[[e] port-numb [detail]]

*Lists the CDP device connected to the specified port. (Allows only one port at a time.) Using **detail** provides a longer list of details on the CDP device the switch detects on the specified port.*

[detail [[e] port-num]]

Provides a list of the details for all of the CDP devices the switch detects. Using port-num produces a list of details for the selected port.

Figure 15-24 lists CDP devices that the switch has detected by receiving their CDP packets.

```
ProCurve> show cdp neighbors
CDP neighbors information
```

Port	Device ID	Platform	Capability
A1	Accounting(0030c1-7fcc40)	J4812A ProCurve Switch...	S
A2	Research(0060b0-889e43)	J4121A ProCurve Switch...	S
A4	Support(0060b0-761a45)	J4121A ProCurve Switch...	S
A7	Marketing(0030c5-38dc59)	J4813A ProCurve Switch...	S
A12	Mgmt NIC(099a05-09df9b)	NIC Model X666	H
A12	Mgmt NIC(099a05-09df11)	NIC Model X666	H

Figure 15-24. Example of CDP Neighbors Table Listing

Enabling CDP Operation. Enabling CDP operation (the default) on the switch causes the switch to add entries to its CDP Neighbors table for any CDP packets it receives from other neighboring CDP devices.

Disabling CDP Operation. Disabling CDP operation clears the switch's CDP Neighbors table and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

Syntax: [no] cdp run

*Enables or disables CDP read-only operation on the switch.
(Default: Enabled)*

For example, to disable read-only CDP on the switch:

```
ProCurve(config)# no cdp run
```

When CDP is disabled:

- **show cdp neighbors** displays an empty CDP Neighbors table
- **show cdp** displays

```
Global CDP information
Enable CDP [Yes]: No
```

Enabling or Disabling CDP Operation on Individual Ports. In the factory-default configuration, the switch has all ports enabled to receive CDP packets. Disabling CDP on a port causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table.

Syntax: [no] cdp enable < [e] port-list >

For example, to disable CDP on port A1:

```
ProCurve(config)# no cdp enable a1
```

—This page is intentionally unused—

File Transfers

Contents

Overview	A-3
Downloading Switch Software	A-3
General Software Download Rules	A-4
Using TFTP To Download Switch Software from a Server	A-4
Menu: TFTP Download from a Server to Primary Flash	A-5
CLI: TFTP Download from a Server to Flash	A-6
Using Secure Copy and SFTP	A-8
How It Works	A-9
The SCP/SFTP Process	A-10
Disable TFTP and Auto-TFTP for Enhanced Security	A-10
Command Options	A-13
Authentication	A-14
SCP/SFTP Operating Notes	A-14
Using Xmodem to Download Switch Software From a PC or UNIX Workstation	A-16
Menu: Xmodem Download to Primary Flash	A-16
CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash	A-17
Switch-to-Switch Download	A-18
Menu: Switch-to-Switch Download to Primary Flash	A-18
CLI: Switch-To-Switch Downloads	A-19
Using PCM+ to Update Switch Software	A-21
Troubleshooting TFTP Downloads	A-21
Transferring Switch Configurations and ACL Command Files	A-23
TFTP: Copying a Configuration from a Remote Host	A-23
TFTP: Copying a Configuration File to a Remote Host	A-24
TFTP: Uploading an ACL Command File from a TFTP Server	A-24
Xmodem: Copying a Configuration File from the Switch to a Serially Connected PC or UNIX Workstation	A-27

Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation	A-27
Copying Diagnostic Data to a Remote Host, PC, or UNIX Workstation	A-29
Copying Command Output to a Destination Device	A-29
Copying Event Log Output to a Destination Device	A-30
Copying Crash Data Content to a Destination Device	A-30
Copying Crash Log Data Content to a Destination Device	A-31

Overview

You can download new switch software, upload or download switch configuration files, and upload command files for configuring Access Control Lists (ACLs).

This appendix includes the following information:

- Downloading switch software (begins below)
- Transferring switch configurations (begins on page A-23)
- Uploading ACL command files (page A-24)

Note that this manual uses the terms *switch software* and *software image* to refer to the downloadable software files of the type the switch uses to operate its networking features. Other terms sometimes used for the same purpose are *Operating System*, or *OS*.

For information on how switch memory operates, including primary and secondary flash, see Chapter 6, “Switch Memory and Configuration”.

Downloading Switch Software

ProCurve periodically provides switch software updates through the ProCurve Networking web site. For more information, refer to the support and warranty booklet shipped with the switch, or visit www.procurve.com and click on **software updates**. After you acquire a new switch software version, you can use one of the following methods for downloading the software to the switch:

Software Download Features

Feature	Default	Menu	CLI	Web
TFTP	n/a	page A-5	page A-6	—
Xmodem	n/a	page A-16	page A-17	—
Switch-to-Switch	n/a	page A-18	page A-19	
Software Update Manager in PCM+	Refer to the documentation provided with PCM+.			

General Software Download Rules

- Switch software that you download via the menu interface always goes to primary flash.
- After a software download, you must reboot the switch to implement the new software. Until a reboot occurs, the switch continues to run on the software it was using before the download commenced.

Note

Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See “Transferring Switch Configurations” on page A-23.

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash. See “Restoring a Flash Image” on page C-58.

Using TFTP To Download Switch Software from a Server

This procedure assumes that:

- A software version for the switch has been stored on a TFTP server accessible to the switch. (The software file is typically available from the ProCurve Networking web site at **www.procurve.com**.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you use the procedure, do the following:

- Obtain the IP address of the TFTP server in which the software file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the software file stored in the TFTP server for the switch (for example, E0820.swi).

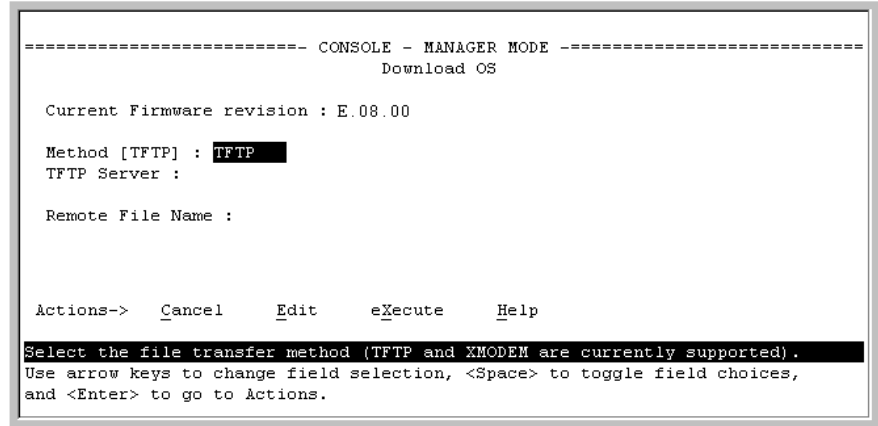
Note

If your TFTP server is a UNIX workstation, *ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.*

Menu: TFTP Download from a Server to Primary Flash

Note that the menu interface accesses only the primary flash.

1. In the console Main Menu, select **Download OS** to display the screen in figure A-A-1. (The term “OS”, or “operating system” refers to the switch software):



```
===== CONSOLE - MANAGER MODE =====
                        Download OS

Current Firmware revision : E.08.00

Method [TFTP] : TFTP
TFTP Server :

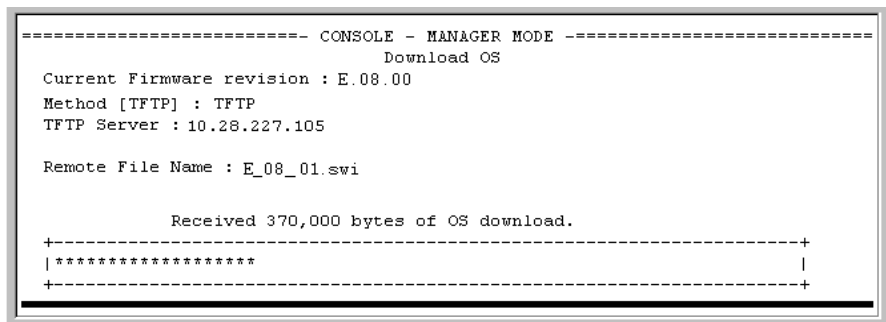
Remote File Name :

Actions->  _Cancel      _Edit      eXecute      _Help

Select the file transfer method (TFTP and XMODEM are currently supported).
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure A-1. Example of a Download OS (Software) Screen (Default Values)

2. Press **[E]** (for **E**dit).
3. Ensure that the **Method** field is set to **TFTP** (the default).
4. In the **TFTP Server** field, type in the IP address of the TFTP server in which the software file has been stored.
5. In the **Remote File Name** field, type the name of the software file. If you are using a UNIX system, remember that the filename is case-sensitive.
6. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download. The following screen then appears:



```
===== CONSOLE - MANAGER MODE =====
                        Download OS

Current Firmware revision : E.08.00
Method [TFTP] : TFTP
TFTP Server : 10.28.227.105

Remote File Name : E_08_01.swi

Received 370,000 bytes of OS download.
+-----+
|*****|
+-----+
```

Figure A-2. Example of the Download OS (Software) Screen During a Download

A “progress” bar indicates the progress of the download. When the entire software file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**

7. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You will then see this prompt:

```
Continue reboot of system? : No
```

Press the space bar once to change No to Yes, then press **[Enter]** to begin the reboot.

Note

When you use the menu interface to download a switch software, the new image is always stored in primary flash. Also, using the Reboot Switch command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI gives you more options. See “Rebooting the Switch” on page 6-18.

8. After you reboot the switch, confirm that the software downloaded correctly:
 - a. From the Main Menu, select **1. Status and Counters**, and from the Status and Counters menu, select **1. General System Information**
 - b. Check the **Firmware revision** line.

CLI: TFTP Download from a Server to Flash

Syntax: copy tftp flash <ip-address> <remote-file> [< primary | secondary >]

This command automatically downloads a switch software file to primary or secondary flash. Note that if you do not specify the flash destination, the Xmodem download defaults to primary flash.

For example, to download a switch software file named E0800.swi from a TFTP server with the IP address of 10.28.227.103 to primary flash:

1. Execute **copy** as shown below:

```
ProCurve# copy tftp flash 10.28.227.103 e0800.swi
The Primary OS Image will be deleted, continue [y/n]? y
01431K
```

Dynamic counter continually displays the number of bytes transferred.

This message means that the image you want to upload will replace the image currently in primary flash.

Figure A-3. Example of the Command to Download an OS (Switch Software)

2. When the switch finishes downloading the software file from the server, it displays this progress message:

Validating and Writing System Software to FLASH ...

3. When the download finishes, you must reboot the switch to implement the newly downloaded software image. To do so, use one of the following commands:

Syntax: boot system flash < primary | secondary >

Boots from the selected flash.

Syntax: reload

Boots from the flash image and startup-config file. On a 5300xl switch running software release E.09.xx or greater (with multiple configuration files), also uses the current startup-config file.

(For more on these commands, refer to “Rebooting the Switch” on page 6-18.)

4. To confirm that the software downloaded correctly, execute **show system** and check the Firmware revision line.

If you need information on primary/secondary flash memory and the boot commands, see “Using Primary and Secondary Flash Image Options” on page 6-13.

Using Secure Copy and SFTP

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session you can then use a third-party software application to take advantage of Secure Copy (SCP) and Secure ftp (SFTP). SCP and SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially you are creating a secure SSH tunnel as a way to transfer files with SFTP and SCP channels.

To use these commands you must install on the administrator workstation a third-party application software client that supports the SFTP and/or SCP functions. Some examples of software that supports SFTP and SCP are PuTTY, Open SSH, WinSCP, and SSH Secure Shell. Most of these are freeware and may be downloaded without cost or licensing from the internet. There are differences in the way these clients work, so be sure you also download the documentation.

As described earlier in this chapter you can use a TFTP client on the administrator workstation to update software images. This is a plain text mechanism and it connects to a standalone TFTP server or another ProCurve switch acting as a TFTP server to obtain the software image file(s). Using SCP and SFTP allows you to maintain your switches with greater security. You can also roll out new software images with automated scripts that make it easier to upgrade multiple switches simultaneously and securely.

SFTP (secure file transfer protocol) is unrelated to FTP, although there are some functional similarities. Once you set up an SFTP session through an SSH tunnel, some of the commands are the same as FTP commands. Certain commands are not allowed by the SFTP server on the switch, such as those that create files or folders. If you try to issue commands such as **create** or **remove** using SFTP the switch server returns an error message.

You can use SFTP just as you would TFTP to transfer files to and from the switch, but with SFTP your file transfers are encrypted and require authentication, so they are more secure than they would be using TFTP. SFTP works only with SSH version 2 (SSH v2).

Note

SFTP over SSH version 1 (SSH v1) is not supported. A request from either the client or the switch (or both) using SSH v1 generates an error message. The actual text of the error message differs, depending on the client software in use. Some examples are:

```
Protocol major versions differ: 2 vs. 1
Connection closed

Protocol major versions differ: 1 vs. 2
Connection closed

Received disconnect from <ip-addr>: /usr/local/
libexec/sftp-server: command not supported
Connection closed
```

SCP (secure copy) is an implementation of the BSD **rcp** (Berkeley UNIX remote copy) command tunneled through an SSH connection.

SCP is used to copy files to and from the switch when security is required. SCP works with both SSH v1 and SSH v2. Be aware that the most third-party software application clients that support SCP use SSHv1.

How It Works

The general process for using SCP and SFTP involves three steps:

1. Open an SSH tunnel between your computer and the switch if you haven't already done so. (This step assumes that you have already set up SSH on the switch.)
2. Execute **ip ssh filetransfer** to tell the switch that you want to enable secure file transfer.
3. Use a third-party client application for SCP and SFTP commands.

The SCP/SFTP Process

To use SCP and SFTP:

1. Open an SSH session as you normally would to establish a secure encrypted tunnel between your computer and the switch. For more detailed directions on how to open an SSH session see the chapter titled “*Configuring Secure Shell (SSH)*” in the *Access Security Guide* for your switch. Please note that this is a one-time procedure for new switches or connections. If you have already done it once you should not need to do it a second time.
2. To enable secure file transfer on the switch (once you have an SSH session established between the switch and your computer), open a terminal window and type in the following command:

```
ProCurve(config)# ip ssh filetransfer
```

Disable TFTP and Auto-TFTP for Enhanced Security

Beginning with software release E.10.02, using the **ip ssh filetransfer** command to enable Secure FTP (SFTP) automatically disables TFTP and auto-TFTP (if either or both are enabled).


```
ProCurve(config)# ip ssh filetransfer
Tftp and auto-tftp have been disabled.
ProCurve(config)# sho run

Running configuration:

; J4850A Configuration Editor; Created on release #E.10.02

hostname "ProCurve"
module 1 type J8161A
module 2 type J8161A
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B24
    ip address 10.28.234.176 255.255.240.0
    exit
ip ssh filetransfer
no tftp-enable
password manager
password operator
```

Enabling SFTP automatically disables TFTP and auto-tftp and displays this message.

Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

Figure A-4. Example of Switch Configuration with SFTP Enabled

If you enable SFTP, then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating rules are:

- The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface, or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

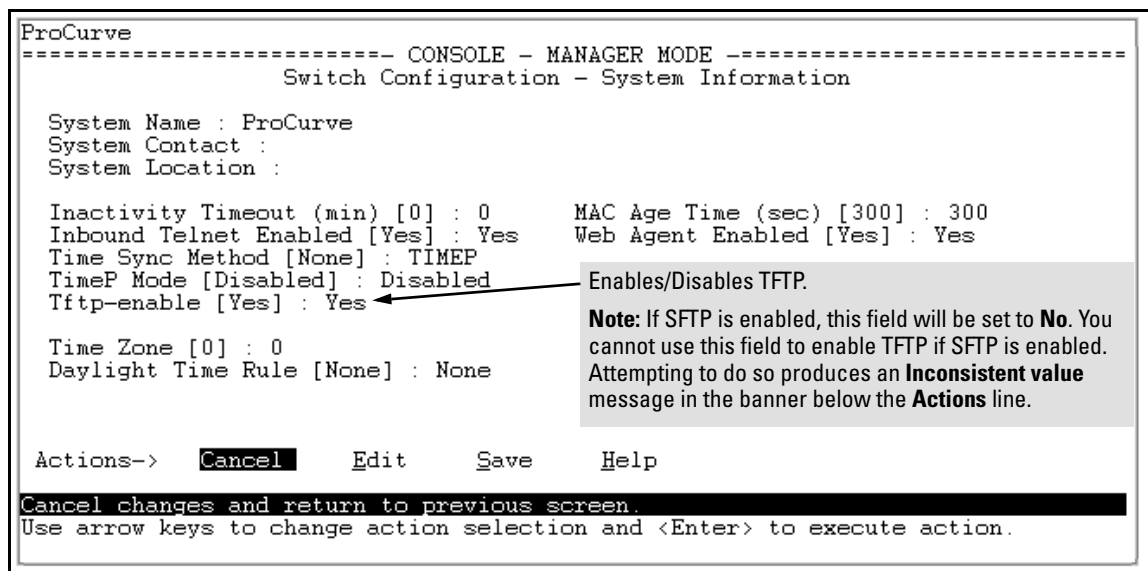


Figure A-5. Using the Menu Interface To Disable TFTP

- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

SFTP must be disabled before enabling tftp.

SFTP must be disabled before enabling auto-tftp.

Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an “inconsistent value” message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but must use the CLI to disable auto-TFTP. The following two CLI commands disable TFTP and auto-TFTP on the switch.

Syntax: no tftp-enable

*This command disables all TFTP operation on the switch except for the auto-TFTP feature. To re-enable TFTP operation, use the **tftp-enable** command. When TFTP is disabled, the instances of **tftp** in the CLI copy command and the Menu interface “Download OS” screen become unavailable.*

Note: This command does **not** disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the **no auto-tftp** command described below to remove the command entry from the switch’s configuration.

Syntax: no auto-tftp

*If auto-TFTP is configured on the switch, this command deletes the **auto-tftp** entry from the switch configuration, thus preventing auto-tftp operation if the switch reboots.*

Note: This command does not affect the current TFTP-enable configuration on the switch.

Command Options

If you need to enable SSH v2 (which is required for SFTP) enter this command:

```
ProCurve(config)# ip ssh version 2
```

Note

As a matter of policy, administrators should *not* enable the SSHv1-only or the SSHv1-or-v2 advertisement modes. SSHv1 is supported on only some legacy switches (such as the ProCurve Series 2500 switches).

To confirm that SSH is enabled type in the command

```
ProCurve(config)# show ip ssh
```

3. Once you have confirmed that you have enabled an SSH session (with the **show ip ssh** command) you can then open your third-party software client application to begin using the SCP or SFTP commands to safely transfer files or issue commands to the switch.

If you need to disable secure file transfer:

```
ProCurve(config)# no ip ssh filetransfer
```

Authentication

Switch memory allows up to ten public keys. This means the authentication and encryption keys you use for your third-party client SCP/SFTP software can differ from the keys you use for the SSH session, even though both SCP and SFTP use a secure SSH tunnel.

Note

SSH authentication through a TACACS+ server and use of SCP or SFTP through an SSH tunnel are mutually exclusive. Thus, if the switch is configured to use TACACS+ for authenticating a secure Telnet SSH session on the switch, you cannot enable SCP or SFTP. Also, if SCP or SFTP is enabled on the switch, you cannot enable TACACS+ authentication for a secure Telnet SSH. On the 5300xl switches, the same mutual exclusion also applies to RADIUS servers. The switch displays a message similar to the following if there is an attempt to configure either option when the other is already configured:

```
RADIUS/TACACS authentication for ssh sessions and
secure file transfer(scp/sftp) may not be configured
simultaneously.
```

To provide username/password authentication on a switch providing SCP or SFTP support, use the switch's local username/password facility. Otherwise, you can use the switch's local public key for authentication.

Some clients such as PSCP (PuTTY SCP) automatically compare switch host keys for you. Other clients require you to manually copy and paste keys to the **\$HOME/.ssh/known_hosts** file. Whatever SCP/SFTP software tool you use, after installing the client software you must verify that the switch host keys are available to the client.

Because the third-party software utilities you may use for SCP/SFTP vary, you should refer to the documentation provided with the utility you select before performing this process.

SCP/SFTP Operating Notes

- When an SFTP client connects, the switch provides a file system displaying all of its available files and folders. No file or directory creation is permitted by the user. Files may only be uploaded or downloaded, according to the permissions mask. All of the necessary files the switch will need are already in place on the switch. You do not need to (nor can you create) new files.

- The switch supports one SFTP session or one SCP session at a time.
- All files have read-write permission. Several SFTP commands, such as create or remove, are not allowed and return an error message. The switch displays the following files:

```

/
+---cfg
|   running-config
|   startup-config
+---log
|   crash-data
|   crash-data-a      (5304xl Only)
|   crash-data-b      "      "
|   crash-data-c      "      "
|   crash-data-d      "      "
|   crash-data-e      (5308xl Only)
|   crash-data-f      "      "
|   crash-data-g      "      "
|   crash-data-h      "      "
|   crash-log
|   crash-log-a       (5304xl Only)
|   crash-log-b       "      "
|   crash-log-c       "      "
|   crash-log-d       "      "
|   crash-log-e       (5308xl Only)
|   crash-log-f       "      "
|   crash-log-g       "      "
|   crash-log-h       "      "
|   event log
+---os
|   primary
|   secondary
\---ssh
    +---mgr_keys
    |   authorized_keys
    \---oper_keys
        authorized_keys
  
```

Once you have configured your switch for secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

Using Xmodem to Download Switch Software From a PC or UNIX Workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port to a PC operating as a terminal. (Refer to the *Installation and Getting Started Guide* you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch software is stored on a disk drive in the PC.
- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the **Send File** option in the **Transfer** dropdown menu.)

Menu: Xmodem Download to Primary Flash

Note that the menu interface accesses only the primary flash.

1. From the console Main Menu, select

7. Download OS

2. Press **[E]** (for **E**dit).
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download. The following message then appears:

**Press enter and then initiate Xmodem transfer
from the attached computer....**

5. Press **[Enter]** and then execute the terminal emulator command(s) to begin Xmodem binary transfer. For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Type the file path and name in the Filename field.
 - c. In the Protocol field, select **Xmodem**.
 - d. Click on the **[Send]** button.

The download will then commence. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.

6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You will then see this prompt:

Continue reboot of system? : No

Press the space bar once to change No to Yes, then press **[Enter]** to begin the reboot.

7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select

1. Status and Counters

1. General System Information

- b. Check the **Firmware revision** line.

CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash

Using Xmodem and a terminal emulator, you can download a software file to either primary or secondary flash.

Syntax: copy xmodem flash [< primary | secondary >]

Downloads a software file to primary or secondary flash. If you do not specify the flash destination, the Xmodem download defaults to primary flash.

For example, to download a switch software file named E0822.swi from a PC (running a terminal emulator program such as HyperTerminal) to primary flash:

1. Execute the following command in the CLI:

```
HPswitch# copy xmodem flash
The Primary OS Image will be deleted, continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Type the file path and name in the Filename field.

- c. In the Protocol field, select **Xmodem**.
- d. Click on the **[Send]** button.

The download can take several minutes, depending on the baud rate used in the transfer.

- 3. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax: boot system flash <primary | secondary>

Reboots from the selected flash.

Syntax: reload

Reboots from the flash image currently in use.

(For more on these commands, see “Rebooting the Switch” on page 6-18.)

- 4. To confirm that the software downloaded correctly:

```
ProCurve> show system
```

Check the **Firmware revision** line. It should show the software version that you downloaded in the preceding steps.

If you need information on primary/secondary flash memory and the boot commands, see “Using Primary and Secondary Flash Image Options” on page 6-13.

Switch-to-Switch Download

You can use TFTP to transfer a software image between two switches of the same series. (For example, all of the Series 5300xl switches use software with the “E” identifier, such as E.08.40.swi, all of the Series 3400cl and Series 6400cl switches use software with the “M” identifier, such as M.08.01, and all of the Series 4200vl switches use the “L” identifier, such as L.10.xx.) The menu interface enables you to transfer primary-to-primary or secondary-to-primary. The CLI enables all combinations of flash location options.

Menu: Switch-to-Switch Download to Primary Flash

Using the menu interface, you can download a switch software file from either the primary or secondary flash of one switch to the primary flash of another switch of the same series.

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default).
3. In the **TFTP Server** field, enter the IP address of the remote switch containing the software file you want to download.
4. For the **Remote File Name**, enter one of the following:
 - To download the software in the primary flash of the source switch, type “**flash**” in lowercase characters.
 - To download the software in the secondary flash of the source switch, type
/os/secondary.
5. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.
6. A “progress” bar indicates the progress of the download. When the entire switch software download has been received, all activity on the switch halts and the following messages appear:

Validating and writing system software to FLASH...

7. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You will then see this prompt:

Continue reboot of system? : No

Press the space bar once to change No to Yes, then press **[Enter]** to begin the reboot.

8. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select

Status and Counters

General System Information

- b. Check the **Firmware revision** line.

CLI: Switch-To-Switch Downloads

Where two switches in your network belong to the same series, you can download a software image between them by initiating a **copy tftp** command from the destination switch. (For example, all of the Series 5300xl switches use software with the “E” identifier, such as E.08.40.swi, all of the Series 3400cl

and Series 6400cl switches use software with the “M” identifier, such as M.08.01 and all of the Series 4200vl switches use the “L” identifier, such as L.10.xx.) The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

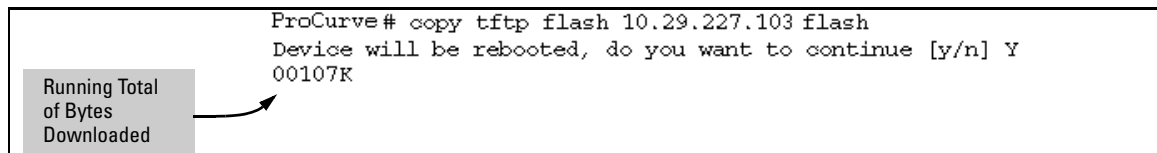
Downloading from Primary Only.

Syntax: copy tftp flash < ip-addr > flash [primary | secondary]

This command (executed in the destination switch) downloads the software flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

For example, to download a software file from primary flash in a with an IP address of 10.29.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:



```
ProCurve# copy tftp flash 10.29.227.103 flash
Device will be rebooted, do you want to continue [y/n] Y
00107K
```

Figure A-6. Switch-To-Switch, from Primary in Source to Either Flash in Destination

Downloading from Either Flash in the Source Switch to Either Flash in the Destination Switch.

Syntax: copy tftp flash < ip-addr > < /os/primary > | < /os/secondary > [primary | secondary]

This command (executed in the destination switch) gives you the most options for downloading between switches. If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

For example, to download a software file from secondary flash in a switch with an IP address of 10.28.227.103 to the secondary flash in a destination switch, you would execute the following command in the destination switch's CLI:

```
ProCurve# copy tftp flash 10.29.227.103 /os/secondary secondary
Device will be rebooted, do you want to continue [y/n] Y
01084K
```

Figure A-7. Switch-to-Switch, from Either Flash in Source to Either Flash in Destination

Using PCM+ to Update Switch Software

ProCurve Manager Plus includes a software update utility for updating on ProCurve switch products such as the 5300xl and 4200vl. (PCM+ version 1.6 and greater will offer this feature for the 3400cl switches and 6400 switches beginning in December, 2004. PCM+ version 2.0.1 will offer this feature for the 4200vl switches.) For further information, refer to the *Getting Started Guide* and the *Administrator's Guide*, provided electronically with the application.

Troubleshooting TFTP Downloads

When using the menu interface, if a TFTP download fails, the Download OS (Operating System, or software) screen indicates the failure.

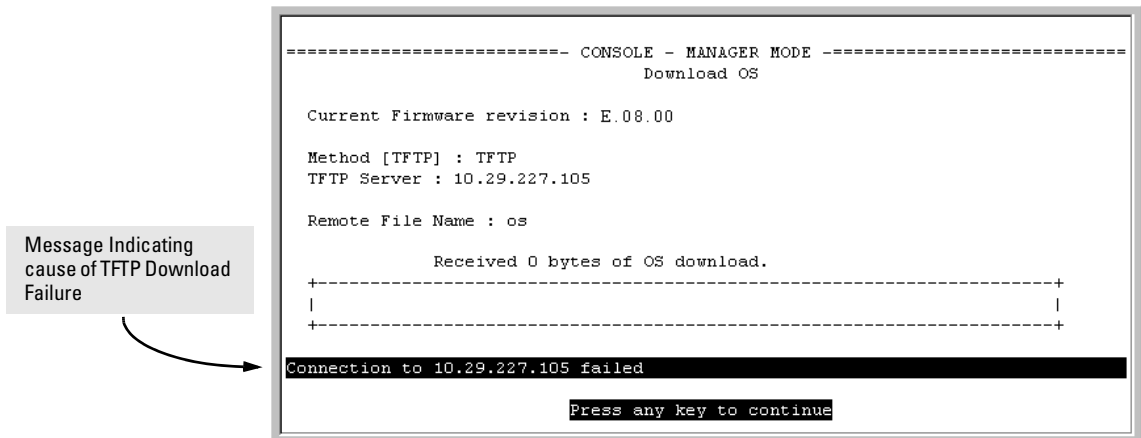


Figure A-8. Example of Message for Download Failure

To find more information on the cause of a download failure, examine the messages in the switch's Event Log by executing this CLI command:

```
ProCurve# show log tftp
```

(For more on the Event Log, see "Using the Event Log To Identify Problem Sources" on "Using the Event Log To Identify Problem Sources" on page C-27.)

Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a UNIX machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the **Download OS** (Operating System, or software) screen.
- One or more of the switch's IP configuration parameters are incorrect.
- For a UNIX TFTP server, the file permissions for the software file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

Note

If an error occurs in which normal switch operation cannot be restored, the switch automatically reboots itself. In this case, an appropriate message is displayed after the switch reboots.

Transferring Switch Configurations and ACL Command Files

Transfer Features

Feature	Default	Menu	CLI	Web
Use TFTP to copy from a remote host to a config file.	n/a	—	below	—
Use TFTP to copy a config file to a remote host.	n/a	—	page A-24	—
Use TFTP to upload and execute a command file for configuring or replacing an ACL in the switch configuration.	n/a	—	page A-24	—
Use Xmodem to copy a configuration from a serially connected host to a config file.	n/a	—	page A-27	—
Use Xmodem to copy a config file to a serially connected host.	n/a	—	page A-27	—

Using the CLI commands described in this section, you can copy switch configurations to and from a switch, or copy an ACL command file to configure or replace an ACL in the switch configuration.

Note

It is useful to note here that you can perform all TFTP operations using SFTP as described in the section on *Using Secure Copy and SFTP* on page A-8 for greater security, if needed.

TFTP: Copying a Configuration from a Remote Host

Syntax: copy tftp < startup-config | running-config > < ip-address > < remote-file > [pc | unix]
copy tftp config < filename > < ip-address > < remote-file > [pc | unix]

All Switches: This command copies a configuration from a remote host to the startup-config or running-config file.

5300xl and 4200vl: On a 5300xl switch running software release E.09.xx or greater or on a 4200vl switch, this command can copy a configuration from a remote host to a designated config file in the switch. For more on multiple configuration files, refer to “Multiple Configuration Files on 5300xl and 4200vl Switches” on page 6-22.

(Refer to “Using Primary and Secondary Flash Image Options” on page 6-13 for more on flash image use.)

For example, to download a configuration file named **sw5300** in the **configs** directory on drive “**d**” in a remote host having an IP address of 10.28.227.105:

```
ProCurve# copy tftp startup-config 10.28.227.105  
                  d:\configs\sw2512
```

TFTP: Copying a Configuration File to a Remote Host

Syntax: copy < startup-config | running-config > tftp < ip-addr > < remote-file > [pc | unix]

copy config < filename > tftp < ip-addr > < remote-file > [pc | unix]

All Switches: This command copies the switch’s startup configuration (startup-config file) or running configuration (running-config file) to a TFTP server.

5300xl or 4200vl: On a 5300xl switch running software release E.09.xx or greater or on a 4200vl switch, this command can copy a designated config file in the switch to a TFTP server. For more on multiple configuration files, refer to “Multiple Configuration Files on 5300xl and 4200vl Switches” on page 6-22.

For example, to upload the current startup configuration to a file named **sw5300** in the configs directory on drive “**d**” in a TFTP server having an IP address of 10.28.227.105:

```
HPswitch# copy startup-config tftp 10.28.227.105  
                  d:\configs\sw5300
```

TFTP: Uploading an ACL Command File from a TFTP Server

This section describes how to upload and execute a command file to the switch for configuring or replacing an Access Control List (ACL) in the switch configuration. Such files should contain only ACE (Access Control Entry) commands. For an example of creating an ACL command file offline, refer to

“Working Offline To Create or Edit an ACL” in the “Access Control Lists (ACLs) chapter of the *Advanced Traffic Management Guide* for your switch.

Syntax: copy tftp command-file < ip-addr > < filename.txt > < unix | pc >

where:

< ip-addr > = The IP address of a TFTP server available to the switch

< filename.txt > = A text file containing ACL commands and stored in the TFTP directory of the server identified by < ip-addr >

< unix | pc > = The type of workstation used for serial, Telnet, or SSH access to the switch CLI

This command copies and executes the named text file from the specified TFTP server address and executes the ACL commands in the file. Depending on the ACL commands used, this action does one of the following in the running-config file:

- Creates a new ACL.
- Replaces an existing ACL. (Refer to “Creating an ACL Offline” in the “Access Control Lists (ACLs)” chapter in the *Advanced Traffic Management Guide* for your switch.)
- Adds to an existing ACL.

For example, suppose you:

1. Created an ACL command file named **vlan10_in.txt** to update an existing ACL.
2. Copied the file to a TFTP server at 18.38.124.16.

Using a PC workstation, you then execute the following from the CLI to upload the file to the switch and implement the ACL commands it contains:

```
ProCurve(config)# copy tftp command-file 18.38.124.16
vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue
[y/n]?
```

File Transfers

Transferring Switch Configurations and ACL Command Files

To continue with the upload, press the **[Y]** key. To abort the upload, press the **[N]** key. Note that if the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice as shown in figure A-9, and continues to implement the remaining ACL commands in the file.

```
ProCurve(config)# copy tftp command-file 18.38.124.16 vlan10_in.txt pc
Running configuration may change, do you want to continue [y/n]? y
  1. ip access-list extended "155"
  2. deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
  3. permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  4. show running
Command files are limited to access-list commands.
  5. exit
ProCurve(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on release #E.08.00

hostname " ProCurve "
cdp run
module 1 type J4820A
ip default-gateway 18.38.248.1
logging 18.38.227.2
snmp-server community "public" Unrestricted
ip access-list extended "155"
  deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
  permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
:
:
:
```

This message indicates that "show running" command just above it is not an ACL command and will be ignored by the switch.

Manually executing **show running** from the CLI indicates that the file was implemented, creating ACL 155 in the switch's running configuration.

Figure A-9. Example of Using the Copy Command to Download and Configure an ACL

For more on this general topic, including an example of an ACL command file created offline, refer to the section titled "Editing ACLs and Creating an ACL Offline" in the "Access Control Lists (ACLs)" chapter of the *Advanced Traffic Management Guide* for your switch.

Xmodem: Copying a Configuration File from the Switch to a Serially Connected PC or UNIX Workstation

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation. You will need to:

- Determine a filename to use.
- Know the directory path you will use to store the configuration file.

Syntax: copy < startup-config | running-config > xmodem < pc | unix >
copy config < filename > xmodem < pc | unix >

All Switches: Uses Xmodem to copy a startup-config or running-config file from the switch to a PC or Unix workstation.

5300xl and 4200vl: A 5300xl switch running software release E.09.xxx or greater, or a 4200vl switch, uses Xmodem to copy a designated configuration file from the switch to a PC or Unix workstation. For more on multiple configuration files, refer to “Multiple Configuration Files on 5300xl and 4200vl Switches” on page 6-22.

For example, to copy a configuration file to a PC serially connected to the switch:

1. Determine the file name and directory location on the PC.
2. Execute the following command:

```
HPswitch# copy startup-config xmodem pc  
Press 'Enter' and start XMODEM on your host...
```

3. After you see the above prompt, press **[Enter]**.
4. Execute the terminal emulator commands to begin the file transfer.

Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation on which is stored the configuration file you want to copy. To complete the copying, you will need to know the name of the file to copy and the drive and directory location of the file.

Syntax: copy xmodem startup-config < pc | unix >
copy xmodem config < filename > < pc | unix >

All Switches: Copies a configuration file from a serially connected PC or UNIX workstation to the switch's startup-config file.

5300xl and 4200vl: 5300xl switches running software release E.09.xx or greater, and 4200vl switches, copy a configuration file from a serially connected PC or UNIX workstation to a designated configuration file on the switch. For more on multiple configuration files, refer to “Multiple Configuration Files on 5300xl and 4200vl Switches” on page 6-22.

For example, to copy a configuration file from a PC serially connected to the switch:

1. Execute the following command:

```
ProCurve# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.
4. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax: boot system flash [primary | secondary]
boot system flash [config < filename >

All Switches: Boots from the selected flash.

5300xl and 4200vl: 5300xl switches running software release E.09.xx or greater, and 4200vl switches, boot from the designated configuration file. For more on multiple configuration files, see “Multiple Configuration Files on 5300xl and 4200vl Switches” on page 6-22.

Syntax: reload

Reboots from the flash image currently in use.

(For more on these commands, see “Rebooting the Switch” on page 6-18.)

Copying Diagnostic Data to a Remote Host, PC, or UNIX Workstation

You can use the CLI to copy the following types of switch data to a text file in a management device:

- **Command Output:** Sends the output of a switch CLI command as a file on the destination device.
- **Event Log:** Copies the switch's Event Log into a file on the destination device.
- **Crash Data:** software-specific data useful for determining the reason for a system crash.
- **Crash Log:** Processor-Specific operating data useful for determining the reason for a system crash.

Copying Command Output to a Destination Device

Syntax: `copy command-output < "cli-command" > tftp < ip-address > < filepath-filename >`

`copy command-output < "cli-command" > xmodem`

These commands direct the displayed output of a CLI command to a file in a destination device.

For example, to use Xmodem to copy the output of **show config** to a serially connected PC:

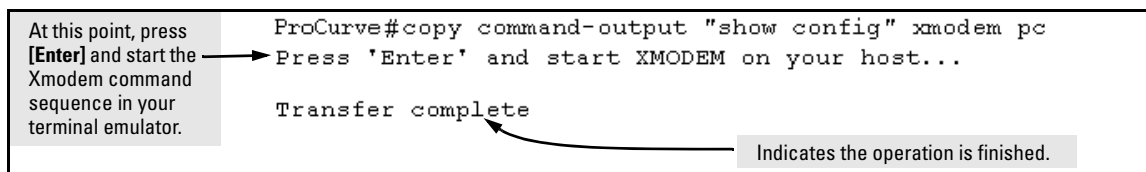


Figure A-10. Example of Sending Command Output to a File on an Attached PC

Note that the command you specify must be enclosed in double-quote marks.

Copying Event Log Output to a Destination Device

Syntax: `copy event-log tftp <ip-address> <filepath_filename>`

`copy event-log xmodem`

These commands use TFTP or Xmodem to copy the Event Log content to a PC or UNIX workstation on the network.

For example, to copy the event log to a PC connected to the switch:

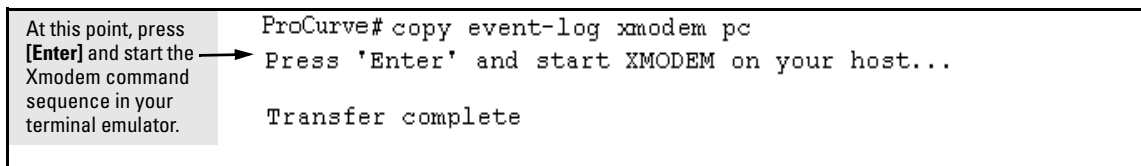


Figure A-11. Example of Sending Event Log Content to a File on an Attached PC

Copying Crash Data Content to a Destination Device

This command uses TFTP or Xmodem to copy the Crash Data content to a PC or UNIX workstation on the network. You can copy individual slot information or the master switch information. If you do not specify either, the command defaults to the master data.

Syntax: `copy crash-data [<slot-id | master>] xmodem`

`copy crash-data [<slot-id | master>] tftp <ip-address> <filename>`

where: slot-id = a - h, and retrieves the crash log or crash data from the processor on the module in the specified slot.

master Retrieves crash log or crash data from the switch's chassis processor.

These commands use TFTP or Xmodem to copy the Event Log content to a PC or UNIX workstation on the network.

For example, to copy the switch's crash data to a file in a PC:

At this point, press **[Enter]** and start the Xmodem command sequence in your terminal emulator.

```
ProCurve(config)# copy crash-data xmodem pc
Press 'Enter' and start XMODEM on your host...
Transfer complete
```

Figure A-12. Example of Copying Switch Crash Data Content to a PC

Copying Crash Log Data Content to a Destination Device

Syntax: copy crash-log [*<slot-id | master>*] tftp *<ip-address>*
<filepath and filename>

```
copy crash-log [<slot-id | master>] xmodem
```

where: slot-id = a-h, and retrieves the crash log or crash data from the processor on the module in the specified slot.

master	<i>Retrieves crash log or crash data from the switch's chassis processor.</i>
--------	---

These commands use TFTP or Xmodem to copy the Crash Log content to a PC or UNIX workstation on the network. You can copy individual slot information or the master switch information. If you do not specify either, the command defaults to the master data.

For example, to copy the Crash Log for slot C to a file in a PC connected to the switch:

At this point, press **[Enter]** and start the Xmodem command sequence in your terminal emulator.

```
ProCurve(config)# copy crash-log c xmodem
Press 'Enter' and start XMODEM on your host...
Transfer complete
```

Figure A-13. Example of sending a Crash Log for Slot C to a File on an Attached PC

File Transfers

Copying Diagnostic Data to a Remote Host, PC, or UNIX Workstation

—This page is intentionally unused—

Monitoring and Analyzing Switch Operation

Contents

Overview	B-3
Status and Counters Data	B-4
Menu Access To Status and Counters	B-5
General System Information	B-5
Menu Access	B-5
CLI Access	B-6
Switch Management Address Information	B-6
Menu Access	B-6
CLI Access	B-7
Module Information	B-8
Menu: Displaying Port Status	B-8
CLI Access	B-8
Port Status	B-9
Menu: Displaying Port Status	B-9
CLI Access	B-9
Web Access	B-9
Viewing Port and Trunk Group Statistics and Flow Control Status	B-10
Menu Access to Port and Trunk Statistics	B-11
CLI Access To Port and Trunk Group Statistics	B-12
Web Browser Access To View Port and Trunk Group Statistics	B-12
Viewing the Switch's MAC Address Tables	B-13
Menu Access to the MAC Address Views and Searches	B-13
CLI Access for MAC Address Views and Searches	B-16
Spanning Tree Protocol (STP) Information	B-17
Menu Access to STP Data	B-17
CLI Access to STP Data	B-18
Internet Group Management Protocol (IGMP) Status	B-19

VLAN Information	B-20
Web Browser Interface Status Information	B-22
Interface Monitoring Features	B-23
Menu: Configuring Port and Static Trunk Monitoring	B-24
CLI: Configuring Port, Mesh, and Static Trunk Monitoring	B-26
Web: Configuring Port Monitoring	B-29

Overview

The switches covered by this guide have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- **Status:** Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data (*page B-4*).
- **Counters:** Display details of traffic volume on individual ports (*page B-10*).
- **Event Log:** Lists switch operating events (“Using the Event Log To Identify Problem Sources” on *page C-27*).
- **Alert Log:** Lists network occurrences detected by the switch—in the Status | Overview screen of the web browser interface (*page 5-16*).
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch. (Refer to “SNMPv1 and SNMPv2c Trap Features” on *page 15-19*.)
- **Port monitoring (mirroring):** Copy all traffic from the specified ports to a designated monitoring port (*page B-23*).

Note

Link test and ping test—analysis tools in troubleshooting situations—are described in appendix C, “Troubleshooting”. See “Diagnostic Tools” on *page C-45*.

Status and Counters Data

This section describes the status and counters screens available through the switch console interface and/or the web browser interface.

Note

You can access all console screens from the web browser interface via Telnet to the console. Telnet access to the switch is available in the Device View window under the **Configuration** tab.

Status or Counters Type	Interface	Purpose	Page
Menu Access to Status and Counters	Menu	Access menu interface for status and counter data.	B-5
General System Information	Menu, CLI	Lists switch-level operating information.	B-5
Management Address Information	Menu, CLI	Lists the MAC address, IP address, and IPX network number for each VLAN or, if no VLANs are configured, for the switch.	B-6
Module Information	Menu, CLI	Lists the module type and description for each slot in which a module is installed.	B-8
Port Status	Menu, CLI, Web	Displays the operational status of each port.	B-9
Port and Trunk Statistics and Flow Control Status	Menu, CLI, Web	Summarizes port activity and lists per-port flow control status.	B-10
VLAN Address Table	Menu, CLI	Lists the MAC addresses of nodes the switch has detected on specific VLANs, with the corresponding switch port.	B-13
Port Address Table	Menu, CLI	Lists the MAC addresses that the switch has learned from the selected port.	B-13
STP Information	Menu, CLI	Lists Spanning Tree Protocol data for the switch and for individual ports. If VLANs are configured, reports on a per-VLAN basis.	B-17
IGMP Status	Menu, CLI	Lists IGMP groups, reports, queries, and port on which querier is located.	B-19
VLAN Information	Menu, CLI	For each VLAN configured in the switch, lists 802.1Q VLAN ID and up/down status.	B-20
Port Status Overview and Port Counters	Web	Shows port utilization and counters, and the Alert Log.	B-22

Menu Access To Status and Counters

Beginning at the Main Menu, display the Status and Counters menu by selecting:

1. Status and Counters

```
===== CONSOLE - MANAGER MODE =====  
Status and Counters Menu  
  
1. General System Information  
2. Switch Management Address Information  
3. Module Information  
4. Port Status  
5. Port Counters  
6. Vlan Address Table  
7. Port Address Table  
8. Spanning Tree Information  
0. Return to Main Menu...  
  
Displays switch management information including software versions.  
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure B-1. The Status and Counters Menu

Each of the above menu items accesses the read-only screens described on the following pages. Refer to the online help for a description of the entries displayed in these screens.

General System Information

Menu Access

From the console Main Menu, select:

1. Status and Counters

1. General System Information

```
=====-- CONSOLE - MANAGER MODE -----  
Status and Counters - General System Information  
  
System Contact      :  
System Location     :  
  
Firmware revision   : E.08.30           Base MAC Addr    : 0001e7-a09900  
ROM Version         : E.05.04           Serial Number     : S2600017409  
  
Up Time             : 2 hours            Memory  - Total   : 24,588,136  
CPU Util (%)        : 1                  Free       : 19,613,568  
  
IP Mgmt  - Pkts Rx : 0                   Packet  - Total   : 832  
          Pkts Tx : 0                   Buffers  Free     : 793  
                                     Lowest    : 769  
                                     Missed     : 0  
                                     24,588,1 6  
  
Actions->  Back  Help  
Return to previous screen.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-2. Example of General Switch Information

This screen dynamically indicates how individual switch resources are being used. See the online Help for details.

CLI Access

Syntax: show system-information

Switch Management Address Information

Menu Access

From the Main Menu, select:

- 1 Status and Counters ...**
- 2. Switch Management Address Information**

```
=====-- CONSOLE - MANAGER MODE -----  
Status and Counters - Management Address Information  
  
Time Server Address : Disabled  
  
VLAN Name      MAC Address      IP Address  
-----  
DEFAULT VLAN   0001e7-a09900    10.28.227.101  
VLAN-22        0001e7-a09900    Disabled  
VLAN-33        0001e7-a09900    Disabled  
  
Actions->      Back      Help  
Return to previous screen.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-3. Example of Management Address Information with VLANs Configured

This screen displays addresses that are important for management of the switch. If multiple VLANs are *not* configured, this screen displays a single IP address for the entire switch. See the online Help for details.

Note

As shown in figure B-3, all VLANs on the switches covered by this guide use the same MAC address. (This includes both the statically configured VLANs and any dynamic VLANs existing on the switch as a result of GVRP operation.)

Also, the switches covered by this guide use a multiple forwarding database. When using multiple VLANs and connecting a switch to a device that uses a single forwarding database, such as a Switch 4000M, there are cabling and tagged port VLAN requirements. For more on this topic, refer to the section titled “Multiple VLAN Considerations” in the “Static Virtual LANs (VLANs)” chapter of the *Advanced Traffic Management Guide* for your switch.

CLI Access

Syntax: show management

Module Information

Use this feature to determine which slots have modules installed and which type(s) of modules are installed.

Menu: Displaying Port Status

From the Main Menu, select:

1. Status and Counters ...
3. Module Information

```
HPswitch
=====
                        Status and Counters - Module Information
=====
Slot   Module Type      Module Description
-----
A      HP J4878A  4x MiniGBIC module
B      HP J4820A  10/100Base-TX module
C      Slot Available
D      Slot Available
E      Slot Available
F      Slot Available
G      Slot Available
H      Slot Available

Actions->  Back      Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-4. Example of Module Information in the Menu Interface

CLI Access

Syntax: show module

Port Status

The web browser interface and the console interface show the same port status data.

Menu: Displaying Port Status

From the Main Menu, select:

- 1. **Status and Counters ...**
- 4. **Port Status**

Status and Counters - Port Status						
Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl
A1		No	Yes	Down		off
A2		No	Yes	Down		off
A3		No	Yes	Down		off
A4		No	Yes	Down		off
B1	10/100TX	No	Yes	Up	100FDx	off
B2	10/100TX	No	Yes	Down	10FDx	off
B3	10/100TX	No	Yes	Down	10FDx	off
B4	10/100TX	No	Yes	Down	10FDx	off
B5	10/100TX	No	Yes	Down	10FDx	off
B6	10/100TX	No	Yes	Down	10FDx	off
B7	10/100TX	No	Yes	Down	10FDx	off
Actions-> Back Intrusion log Help						
Return to previous screen.						
Use up/down arrow keys to scroll to other entries, left/right arrow keys to change action selection, and <Enter> to execute action.						

Figure B-5. Example of Port Status on the Menu Interface

CLI Access

Syntax: show interfaces brief

Web Access

- 1. Click on the **Status** tab.
- 2. Click on **[Port Status]**.

Viewing Port and Trunk Group Statistics and Flow Control Status

Feature	Default	Menu	CLI	Web
viewing port and trunk statistics for all ports, and flow control status	n/a	page B-11	page B-12	page B-12
viewing a detailed summary for a particular port or trunk	n/a	page B-11	page B-12	page B-12
resetting counters	n/a	page B-11	page B-12	page B-12

These features enable you to determine the traffic patterns for each port since the last reboot or reset of the switch. You can display:

- A general report of traffic on all LAN ports and trunk groups in the switch, along with the per-port flow control status (On or Off).
- A detailed summary of traffic on a selected port or trunk group.

You can also reset the counters for a specific port.

The menu interface and the web browser interface provide a dynamic display of counters summarizing the traffic on each port. The CLI lets you see a static “snapshot” of port or trunk group statistics at a particular moment.

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting. See the “Note On Reset”, below.

Note on Reset

The **Reset** action resets the counter display to zero for the current session, but does not affect the cumulative values in the actual hardware counters. (In compliance with the SNMP standard, the values in the hardware counters are not reset to zero unless you reboot the switch.) Thus, using the **Reset** action resets the displayed counters to zero for the current session only. Exiting from the console session and starting a new session restores the counter displays to the accumulated values in the hardware counters.

Menu Access to Port and Trunk Statistics

To access this screen from the Main Menu, select:

1. Status and Counters ...

4. Port Counters

```

=====  CONSOLE - MANAGER MODE  =====
                        Status and Counters - Port Counters

```


Port	Total Bytes	Total Frames	Errors Rx	Drops Tx	Flow Ctrl
A1	195,072	323	0	0	off
A2	651,816	871	0	0	off
A3-Trk1	290,163	500	0	0	off
A4-Trk1	260,134	501	0	0	off
C1	859,363	5147	0	0	off
C2	674,574	1693	0	0	off
C3	26,554	246	0	0	off
C4	113,184	276	0	0	off
C5	0	0	0	0	off

```

Actions->  Back  Show details  Reset  Help
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.

```

Figure B-6. Example of Port Counters on the Menu Interface

To view details about the traffic on a particular port, use the  key to highlight that port number, then select **Show Details**. For example, selecting port A2 displays a screen similar to figure B-7, below.

```

=====  CONSOLE - MANAGER MODE  =====
                        Status and Counters - Port Counters - Port A2

```

Link Status	: up		
Bytes Rx	: 630,746	Bytes Tx	: 21,070
Unicast Rx	: 568	Unicast Tx	: 285
Bcast/Mcast Rx	: 18	Bcast/Mcast Tx	: 0
FCS Rx	: 0	Drops Tx	: 0
Alignment Rx	: 0	Collisions Tx	: 0
Runts Rx	: 0	Late Colln Tx	: 0
Giants Rx	: 0	Excessive Colln	: 0
Total Rx Errors	: 0	Deferred Tx	: 0

```

Actions->  Back  Reset  Help
Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.

```

Figure B-7. Example of the Display for Show details on a Selected Port

This screen also includes the **Reset** action for the current session. (See the “Note on Reset” on page B-10.)

CLI Access To Port and Trunk Group Statistics

To Display the Port Counter Summary Report.

Syntax: show interfaces

This command provides an overview of port activity for all ports on the switch.

To Display a Detailed Traffic Summary for Specific Ports. .

Syntax: show interfaces < port-list >

This command provides traffic details for the port(s) you specify

To Reset the Port Counters for a Specific Port.

Syntax: clear statistics < port-list >

This command resets the counters for the specified ports to zero for the current session. (See the “Note on Reset” on page B-10.)

Web Browser Access To View Port and Trunk Group Statistics

1. Click on the **Status** tab.
2. Click on **[Port Counters]**.
3. To refresh the counters for a specific port, click anywhere in the row for that port, then click on **[Refresh]**.

Note

The reset the port counters to zero, you must reboot the switch.

Viewing the Switch’s MAC Address Tables

Feature	Default	Menu	CLI	Web
viewing MAC addresses on all ports on a specific VLAN	n/a	page B-13	page B-16	—
viewing MAC addresses on a specific port	n/a	page B-15	page B-16	—
searching for a MAC address	n/a	page B-15	page B-16	—

These features help you to view:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

Menu Access to the MAC Address Views and Searches

Per-VLAN MAC-Address Viewing and Searching. This feature lets you determine which switch port on a selected VLAN is being used to communicate with a specific device on the network. The per-VLAN listing includes:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

1. From the Main Menu, select:

1. Status and Counters
5. VLAN Address Table

2. The switch then prompts you to select a VLAN.



3. Use the Space bar to select the VLAN you want, then press **[Enter]**. The switch then displays the MAC address table for that VLAN:

```
----- CONSOLE - MANAGER MODE -----
Status and Counters - Address Table

  MAC Address  Located on Port
-----
0030c1-7f49c0  A3
0030c1-7fec40  A1
0030c1-b29ac0  A3
0060b0-17de5b  A3
0060b0-880a80  A2
0060b0-df1a00  A3
0060b0-df2a00  A3
0060b0-e9a200  A3
009027-e74f90  A3
080009-21ae84  A3
080009-62c411  A3
080009-6563e2  A3

Actions->  Back  Search  Next page  Prev page  Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-8. Example of the Address Table

To page through the listing, use **Next page** and **Prev page**.

Finding the Port Connection for a Specific Device on a VLAN. This feature uses a device's MAC address that you enter to identify the port used by that device.

1. Proceeding from figure B-8, press [**S**] (for **Search**), to display the following prompt:
Enter MAC address: _
2. Type the MAC address you want to locate and press [**Enter**]. The address and port number are highlighted if found. If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Located MAC Address and Corresponding Port Number	<pre>----- CONSOLE - MANAGER MODE ----- Status and Counters - Address Table MAC Address Located on Port ----- 0030c1-7fcc6d 2 005004-17df9c 1 0060b0-889e00 1</pre>
---	--

Figure B-9. Example of Menu Indicating Located MAC Address

3. Press [**P**] (for **Prev page**) to return to the full address table listing.

Port-Level MAC Address Viewing and Searching. This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

1. From the Main Menu, select:

- 1. Status and Counters**
7. Port Address Table

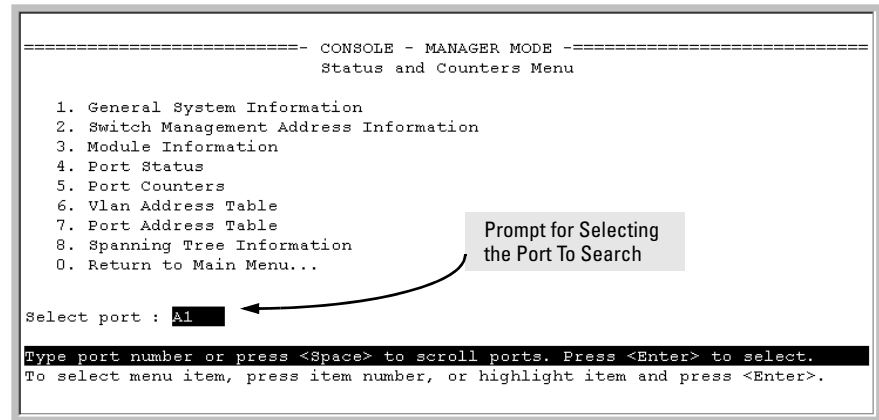


Figure B-10. Listing MAC Addresses for a Specific Port

2. Use the Space bar to select the port you want to list or search for MAC addresses, then press **[Enter]** to list the MAC addresses detected on that port.

Determining Whether a Specific Device Is Connected to the Selected Port. Proceeding from step 2, above:

1. Press **[S]** (for **Search**), to display the following prompt:
Enter MAC address: _
2. Type the MAC address you want to locate and press **[Enter]**. The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.
3. Press **[P]** (for **Prev page**) to return to the previous per-port listing.

CLI Access for MAC Address Views and Searches

Syntax: show mac-address
 [vlan < *vlan-id* >]
 [< *port-list* >]
 [< mac-addr >]

To List All Learned MAC Addresses on the Switch, with The Port Number on Which Each MAC Address Was Learned.

```
ProCurve> show mac-address
```

To List All Learned MAC Addresses on one or more ports, with Their Corresponding Port Numbers. For example, to list the learned MAC address on ports A1 through A4 and port A6:

```
ProCurve> show mac-address a1-a4,a6
```

To List All Learned MAC Addresses on a VLAN, with Their Port Numbers. This command lists the MAC addresses associated with the ports for a given VLAN. For example:

```
ProCurve> show mac-address vlan 100
```

Note

The switches covered by this guide operate with a multiple forwarding database architecture. For more on this topic, refer to “Duplicate MAC Addresses on Different Switches” on page C-14

To Find the Port On Which the Switch Learned a Specific MAC Address. For example, to find the port on which the switch learns a MAC address of 080009-21ae84:

```
ProCurve# show mac-address 080009-21ae84
Status and Counters - Address Table - 080009-21ae84
MAC Address : 080009-21ae84
Located on Port : A2
```

Spanning Tree Protocol (STP) Information

Menu Access to STP Data

From the Main Menu, select:

1. Status and Counters ...
8. Spanning Tree Information

STP must be enabled on the switch to display the following data:

```
----- CONSOLE - MANAGER MODE -----
                        Status and Counters - Spanning Tree Information

STP Enabled           : Yes
Switch Priority        : 32,768
Hello Time            : 2
Max Age               : 20
Forward Delay         : 15

Topology Change Count : 3
Time Since Last Change : 4 mins

Root MAC Address      : 0030c1-7fcc40
Root Path Cost        : 0
Root Port             : This switch is root
Root Priority          : 32768

Actions->  Back      Show ports  Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-11. Example of Spanning Tree Information

Use this screen to determine current switch-level STP parameter settings and statistics.

You can use the **Show ports** action at the bottom of the screen to display port-level information and parameter settings for each port in the switch (including port type, cost, priority, operating state, and designated bridge) as shown in figure B-12.

```
===== CONSOLE - MANAGER MODE =====
Status and Counters - Spanning Tree - Port Information
```

Port	Type	Cost	Priority	State	Designated Bridge
A1	100/1000T	5	128	Forwarding	0001e7-a09900
A2	100/1000T	5	128	Forwarding	0001e7-a09900
A3	100/1000T	5	128	Disabled	
A4	100/1000T	5	128	Disabled	
A5	100/1000T	5	128	Disabled	
A6	100/1000T	5	128	Disabled	
C1	1000SX	5	128	Forwarding	0001e7-a09900
C2	1000SX	5	128	Forwarding	0001e7-a09900
C3	1000SX	5	128	Forwarding	0001e7-a09900

```
Actions-> Back Help
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-12. Example of STP Port Information

CLI Access to STP Data

This option lists the STP configuration, root data, and per-port data (cost, priority, state, and designated bridge).

Syntax: show spanning-tree

ProCurve> show spanning-tree

Internet Group Management Protocol (IGMP) Status

The switch uses the CLI to display the following IGMP status on a per-VLAN basis:

Show Command	Output
show ip igmp	Global command listing IGMP status for all VLANs configured in the switch: <ul style="list-style-type: none">• VLAN ID (VID) and name• Active group addresses per VLAN• Number of report and query packets per group• Querier access port per VLAN
show ip igmp <vlan-id>	Per-VLAN command listing above IGMP status for specified VLAN (VID)
show ip igmp group <ip-addr>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.

For example, suppose that **show ip igmp** listed an IGMP group address of 224.0.1.22. You could get additional data on that group by executing the following:

```
HPswitch> show ip igmp group 224.0.1.22

IGMP ports for group 224.0.1.22

  Port Type      Access      Age Timer  Leave Timer
  ---- -
  3    10/100TX  host        0          0
```

Figure B-13. Example of IGMP Group Data

VLAN Information

The switch uses the CLI to display the following VLAN status:

Show Command	Output
show vlan	Lists: <ul style="list-style-type: none">• Maximum number of VLANs to support• Existing VLANs• Status (static or dynamic)• Primary VLAN
show vlan <vlan-id>	For the specified VLAN, lists: <ul style="list-style-type: none">• Name, VID, and status (static/dynamic)• Per-Port mode (tagged, untagged, forbid, no/auto)• “Unknown VLAN” setting (Learn, Block, Disable)• Port status (up/down)

For example, suppose that your switch has the following VLANs:

Ports	VLAN	VID
A1 - A12	DEFAULT_VLAN	1
A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

The next three figures show how you could list data on the above VLANs.

Listing the VLAN ID (VID) and Status for ALL VLANs in the Switch.

ProCurve> show vlan			
Status and Counters - VLAN Information			
VLAN support : Yes			
Maximum VLANs to support : 9			
Primary VLAN: DEFAULT_VLAN			
802.1Q	VLAN ID	Name	Status
-----	-----	-----	----
	1	DEFAULT_VLAN	Static
	33	VLAN-33	Static
	44	VLAN-44	Static

Figure B-14. Example of VLAN Listing for the Entire Switch

Listing the VLAN ID (VID) and Status for Specific Ports.

Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

```
ProCurve>show vlan ports A1-A2
Status and Counters - VLAN Information - for ports A1,A2
802.1Q VLAN ID Name          Status
-----
1          DEFAULT_VLAN      Static
33         VLAN-33           Static
```

Figure B-15. Example of VLAN Listing for Specific Ports

Listing Individual VLAN Status.

```
ProCurve>show vlan 1
Status and Counters - VLAN Information - Ports - VLAN 1
802.1Q VLAN ID : 1
Name           : DEFAULT_VLAN
Status        : Static

Port Information Mode      Unknown VLAN Status
-----
A1              Untagged Learn      Up
A2              Tagged   Learn      Up
A3              Untagged Learn      Up
A4              Untagged Learn      Down
A5              Untagged Learn      Down
*              *              *
*              *              *
*              *              *
```

Figure B-16. Example of Port Listing for an Individual VLAN

Web Browser Interface Status Information

The “home” screen for the web browser interface is the Status Overview screen, as shown below. As the title implies, it provides an overview of the status of the switch, including summary graphs indicating the network utilization on each of the switch ports, symbolic port status indicators, and the Alert Log, which informs you of any problems that may have occurred on the switch.

For more information on this screen, refer to chapter 5, “Using the Web Browser Interface” .

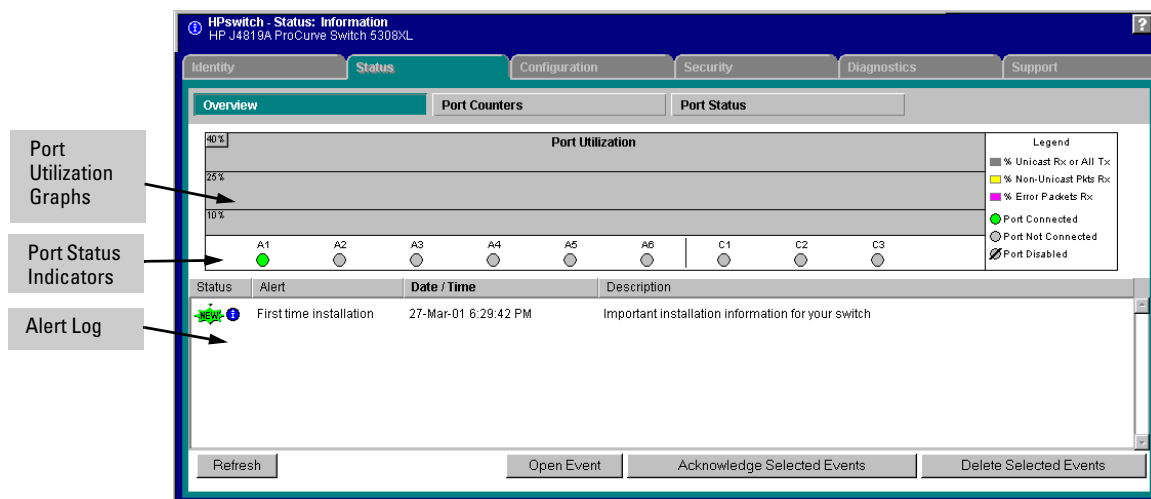


Figure B-17. Example of a Web Browser Interface Status Overview Screen

Interface Monitoring Features

Port Monitoring Features

Feature	Default	Menu	CLI	Web
display monitoring configuration	disabled	page B-24	page B-26	page B-29
configure the monitor port(s)	ports: none	page B-24	page B-27	page B-29
selecting or removing ports	none selected	page B-24	page B-28	page B-29

You can designate monitoring of inbound and outbound traffic on:

- **Ports and static trunks:** Allows monitoring of individual ports, groups of contiguous ports, and static port trunks.
- **Meshed ports:** Allows traffic monitoring on all ports configured for meshing on the switch.
- **Static VLANs:** Allows traffic monitoring on one static VLAN (*5300xl switches and 4200vl switches only*).

The switch monitors network activity by copying all traffic inbound and outbound on the specified interfaces to the designated monitoring port, to which a network analyzer can be attached.

Note

VLANs, a switch mesh, and port trunks cannot be used as a monitoring port.

The switch can monitor static LACP trunks, but not dynamic LACP trunks.

It is possible, when monitoring multiple interfaces in networks with high traffic levels, to copy more traffic to a monitor port than the link can support. In this case, some packets may not be copied to the monitor port.

Menu: Configuring Port and Static Trunk Monitoring

This procedure describes configuring the switch for monitoring when monitoring is disabled. (If monitoring has already been enabled, the screens will appear differently than shown in this procedure.)

1. From the Console Main Menu, Select:

2. Switch Configuration...

3. Network Monitoring Port

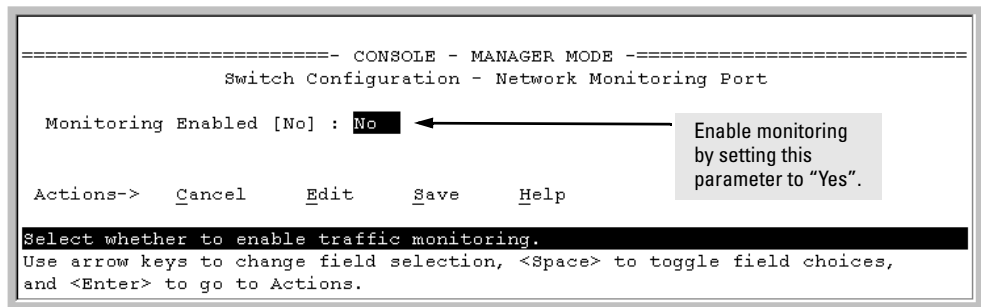


Figure B-18. The Default Network Monitoring Configuration Screen

2. In the Actions menu, press [E] (for Edit).
3. If monitoring is currently disabled (the default) then enable it by pressing the Space bar (or [Y]) to select Yes.
4. Press the down arrow key to display a screen similar to the following and move the cursor to the **Monitoring Port** parameter.

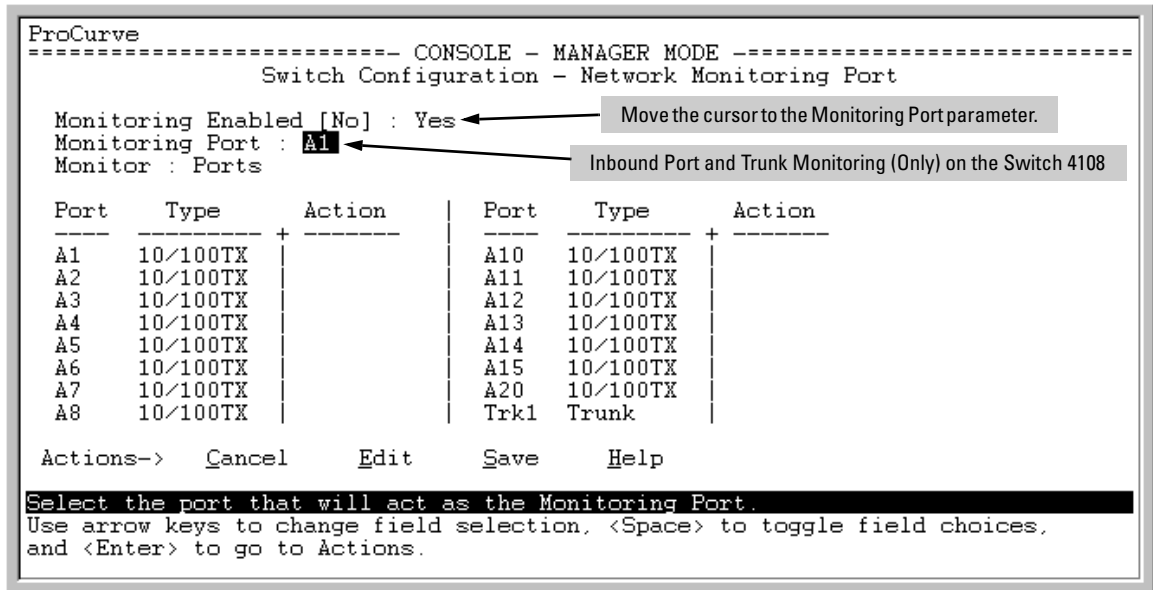


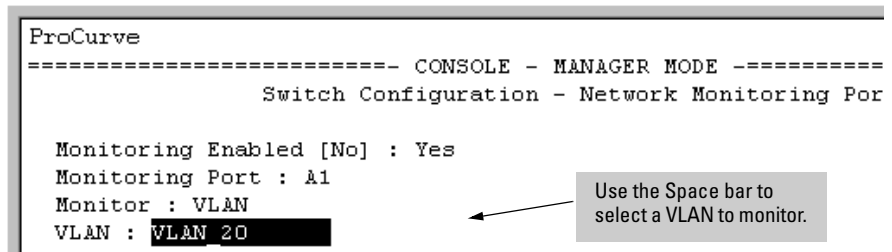
Figure B-19. How To Select a Monitoring Port

5. Use the Space bar to select the port to use for monitoring.
6. Highlight the Monitor field and use the Space bar to select the interfaces to monitor:

Ports: Use for monitoring ports, static trunks, or the mesh.

VLAN: Use for monitoring a VLAN (*5300xl and 4200vl switches*).

7. Do one of the following:
 - If you are monitoring ports, static trunks, or the mesh, go to step 8.
 - If you are monitoring a VLAN on a 5300xl switch:
 - i. Press **[Tab]** or the down arrow key to move to the **VLAN** field.



- ii. Use the Space bar to select the VLAN you want to monitor.
- iii. Go to step 10.
8. Use the down arrow key to move the cursor to the **Action** column for the individual ports and position the cursor at a port you want to monitor.
9. Press the Space bar to select **Monitor** for each port and trunk that you want monitored. (Use the down arrow key to move from one interface to the next in the **Action** column.)
10. When you finish selecting ports to monitor, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
11. Return to the Main Menu.

CLI: Configuring Port, Mesh, and Static Trunk Monitoring

Port, Mesh, and Static Trunk Monitoring Commands Used in This Section

show monitor	below
mirror-port	page B-27
monitor	page B-28

You must use the following configuration sequence to configure port and static trunk monitoring in the CLI:

1. Assign a monitoring (mirror) port.
2. Designate the port(s), mesh, and/or static trunk(s) to monitor.

Displaying the Monitoring Configuration.

Syntax: show monitor

This command lists the port assigned to receive monitored traffic and the ports and/or trunks being monitored.

For example, if you assign port A6 as the monitoring port and configure the switch to monitor ports A1 - A3, **show monitor** displays the following:

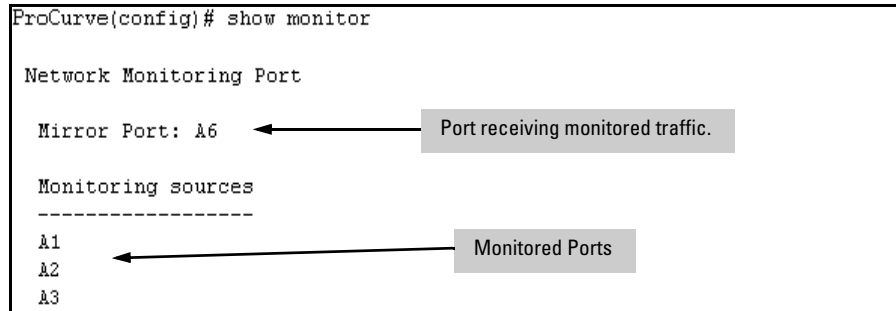


Figure B-20. Example of Monitored Port Listing

Configuring the Monitor Port.

Syntax: [no] mirror-port [< port-num >]

This command assigns or removes a monitoring port, and must be executed from the global configuration level. Removing the monitor port disables port monitoring and resets the monitoring parameters to their factory-default settings.

For example, to assign port A6 as the monitoring port:

```
ProCurve(config)# mirror-port a6
```

To turn off monitoring:

```
ProCurve(config)# no mirror-port
```

Selecting or Removing Monitoring Source Interfaces. After you configure a monitor port you can use either the global configuration level or the interface context level to select ports, static trunks, meshed ports, or (for the 5300xl switches or 4200vl switches) VLANs as monitoring sources. You can also use either level to remove monitoring sources.

Syntax: [no] interface < monitor-list > monitor
[no] vlan < vid > monitor

where:

< monitor-list > Includes port numbers, static trunk names, and meshing, such as **a4**, **c7**, **b5-b8**, **trk1**, and **mesh**.

< vid > Allows monitoring of one VLAN. (This option applies to the 5300xl and 4200vl switches.)

Identifies the switch elements to monitor through the currently configured monitor port. You can monitor the port(s), static trunk(s), and any switch mesh available on the switch or, on a 5300xl or 4200vl, one VLAN.

Note

Individual ports, static trunks, and meshing can all be monitored at the same time. However, if you configure the switch to monitor a VLAN (5300xl and 4200vl switches only), all other interfaces are removed from monitoring. Also, you can configure only one VLAN at a time for monitoring.

Elements in the monitor list can include port numbers, static trunk names, and the mesh at the same time.

For example, with a port such as port A6 configured as the monitoring (mirror) port, you would use either of the following commands to select these interfaces for monitoring:

- A1 through A3, and A5
- Trunks 1 and 2
- Meshing

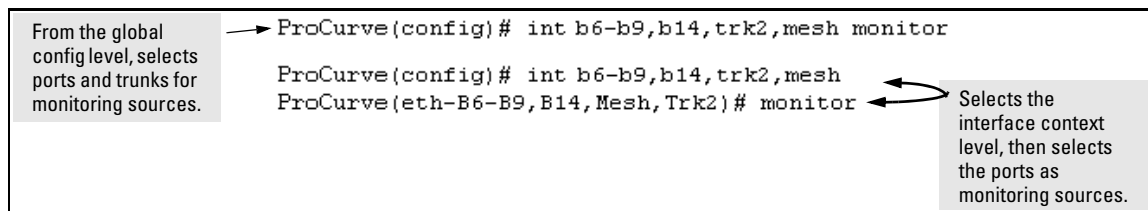


Figure B-21. Examples of Selecting Ports and Static Trunks as Monitoring Sources

To monitor a VLAN:

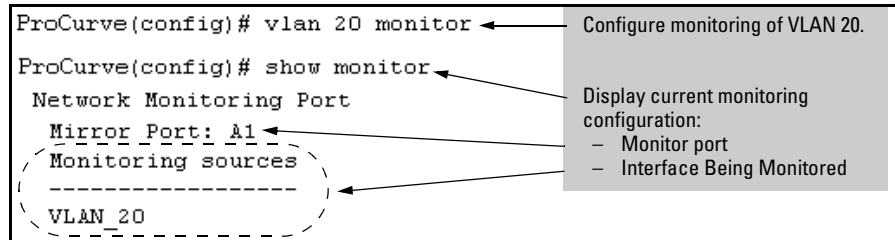


Figure B-22. Example of Configuring VLAN Monitoring

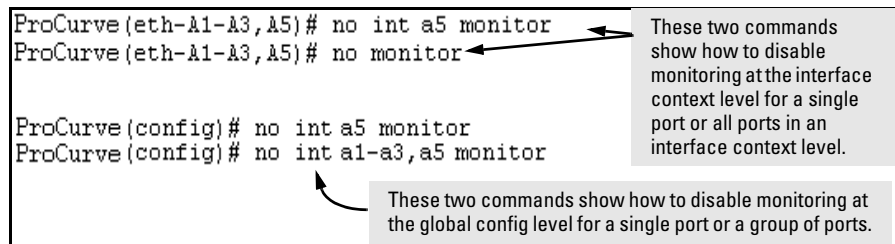


Figure B-23. Examples of Removing Ports as Monitoring Sources

Web: Configuring Port Monitoring

To enable port monitoring:

1. Click on the **Configuration** tab.
2. Click on **[Monitor Port]**.
3. To monitor one or more ports.
 - a. Click on the radio button for **Monitor Selected Ports**.
 - b. Select the port(s) to monitor.
4. Click on **[Apply Changes]**.

To remove port monitoring:

1. Click on the **[Monitoring Off]** radio button.
2. Click on **[Apply Changes]**.

For web-based Help on how to use the web browser interface screen, click on the **[?]** button provided on the web browser screen.

— This page is intentionally unused. —

Troubleshooting

Contents

Overview	C-3
Troubleshooting Approaches	C-4
Browser or Telnet Access Problems	C-5
Unusual Network Activity	C-7
General Problems	C-7
802.1Q Prioritization Problems	C-8
ACL Problems	C-8
IGMP-Related Problems	C-13
LACP-Related Problems	C-13
Mesh-Related Problems	C-14
Port-Based Access Control (802.1x)-Related Problems	C-15
QoS-Related Problems	C-18
Radius-Related Problems	C-18
Spanning-Tree Protocol (STP) and Fast-Uplink Problems	C-19
SSH-Related Problems	C-20
TACACS-Related Problems	C-22
TimeP, SNTP, or Gateway Problems	C-24
VLAN-Related Problems	C-24
Using the Event Log To Identify Problem Sources	C-27
Menu: Entering and Navigating in the Event Log	C-29
CLI: Listing Events	C-30
Reducing Duplicate Event Log and SNMP Trap Messages	C-31
Debug and Syslog Messaging Operation	C-34
Debug Command Operation	C-35

Debug Types	C-36
Debug Destinations	C-38
Syslog Operation	C-39
Viewing the Debug Configuration	C-40
Steps for Configuring Debug and Syslog Messaging	C-40
Operating Notes for Debug and Syslog	C-44
Diagnostic Tools	C-45
Port Auto-Negotiation	C-45
Ping and Link Tests	C-45
Web: Executing Ping or Link Tests	C-47
CLI: Ping or Link Tests	C-48
Displaying the Configuration File	C-50
CLI: Viewing the Configuration File	C-50
Web: Viewing the Configuration File	C-50
Listing Switch Configuration and Operation Details	C-50
CLI Administrative and Troubleshooting Commands	C-52
Traceroute Command	C-53
Restoring the Factory-Default Configuration	C-57
CLI: Resetting to the Factory-Default Configuration	C-57
Clear/Reset: Resetting to the Factory-Default Configuration ..	C-57
Restoring a Flash Image	C-58

Overview

This chapter addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, refer to the installation guide you received with the switch.)

Note

ProCurve periodically places switch software updates on the ProCurve Networking web site. ProCurve recommends that you check this web site for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, see the Support and Warranty booklet shipped with the switch.

Troubleshooting Approaches

Use these approaches to diagnose switch problems:

- Check the ProCurve Networking web site for software updates that may have solved your problem: **www.procurve.com**
- Check the switch LEDs for indications of proper switch operation:
 - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
 - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.

See the *Installation Guide* shipped with the switch for a description of the LED behavior and information on using the LEDs for troubleshooting.

- Check the network topology/installation. See the *Installation Guide* shipped with the switch for topology information.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. See the *Installation Guide* shipped with the switch for correct cable types and connector pin-outs.
- Use ProCurve Manager to help isolate problems and recommend solutions.
- Use the Port Utilization Graph and Alert Log in the web browser interface included in the switch to help isolate problems. See Chapter 5, “Using the Web Browser Interface” for operating information. These tools are available through the web browser interface:
 - Port Utilization Graph
 - Alert Log
 - Port Status and Port Counters screens
 - Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. See chapter 4, “Using the Switch Console Interface” for operating information. These tools are available through the switch console
 - Status and Counters screens
 - Event Log
 - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet Access Problems

Cannot access the web browser interface:

- Access may be disabled by the **Web Agent Enabled** parameter in the switch console. Check the setting on this parameter by selecting:

2. Switch Configuration ...

1. System Information

- The switch may not have the correct IP address, subnet mask or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration ...

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

1. Status and Counters ...

2. Switch Management Address Information

also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows web browser access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch web browser interface to operate correctly. See the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network:

- Off subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the `ip route` command to configure a static (default) route before enabling routing. Refer to chapter 16, “IP Routing Features”, for more information.
- Telnet access may be disabled by the **Inbound Telnet Enabled** parameter in the System Information screen of the menu interface:

2. Switch Configuration

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch’s Console port and selecting:

2. Switch Configuration

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, see the **Note**, above.

- If you are using DHCP to acquire the IP address for the switch, the IP address “lease time” may have expired so that the IP address has changed. For more information on how to “reserve” an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.

Unusual Network Activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switch console interface or with a network management tool such as ProCurve Manager. Refer to the *Installation Guide* you received with the switch for information on using LEDs to identify unusual network activity.

A topology loop can also cause excessive network activity. The Event Log “FFI” messages can be indicative of this type of problem.

General Problems

The network runs slow; processes fail; users cannot access servers or other devices. Broadcast storms may be occurring in the network. These may be due to redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (i.e. topology loops) that will cause broadcast storms.
- Turn on Spanning Tree Protocol to block redundant links (i.e. topology loops)
- Check for FFI messages in the Event Log.

Duplicate IP Addresses. This is indicated by this Event Log message:

ip: Invalid ARP source: IP address on IP address

where: both instances of *IP address* are the same address, indicating the switch’s IP address has been duplicated somewhere on the network.

Duplicate IP Addresses in a DHCP Network. If you use a DHCP server to assign IP addresses in your network and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server “leases” the address to another device.

This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure “reservations” in the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, refer to the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source: < IP-address > on <IP-address>
```

where: both instances of *IP-address* are the same address, indicating the IP address that has been duplicated somewhere on the network.

The Switch Has Been Configured for DHCP/Bootp Operation, But Has Not Received a DHCP or Bootp Reply. When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1Q Prioritization Problems

Ports configured for non-default prioritization (level 1 - 7) are not performing the specified action. If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

ACL Problems

Series 5300xl Switches Only: ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.

1. The 5300xl switch may be running with IP routing disabled. To ensure that IP routing is enabled, execute **show running** and look for the IP routing statement in the resulting listing. For example:

```

ProCurve(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on release #E.08.01

hostname " ProCurve "
cdp run
module 1 type J4820A
ip default-gateway 10.30.248.1
ip routing
logging 10.28.227.2
snmp-server community "public" Unrestricted
ip access-list extended "Controls for VLAN 20"
  permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
  permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
  deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
  deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
  deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
  permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
exit
-- MORE --, next page: Space, next line: Enter, quit: Control-C

```

Indicates that routing is enabled; a requirement for ACL operation. (There is an exception. See the **Note**, below.)

Figure C-1. Indication that Routing Is Enabled

Note

If an ACL assigned to a VLAN includes an ACE referencing an IP address on the switch itself as a packet source or destination, the ACE screens traffic to or from this switch address regardless of whether IP routing is enabled. This is a security measure designed to help protect the switch from unauthorized management access.

If you need to configure IP routing, execute the **ip routing** command.

2. ACL filtering on the 5300xl switches applies only to routed packets and packets having a destination IP address (DA) on the switch itself. Also, the switch applies assigned ACLs only at the point where traffic enters or leaves the switch on a VLAN. Ensure that you have correctly applied your ACLs ("in" and/or "out") to the appropriate VLAN(s).

The switch does not allow management access from a device on the same VLAN.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL always blocks packets having the same DA as the switch's IP address on the same VLAN. That is, bridged packets with the switch itself as the destination are blocked as a security measure. To preempt this action, edit the ACL to include an ACE that permits access to the switch's DA on that VLAN from the management device.

Error (Invalid input) when entering an IP address.

When using the “host” option in the command syntax, ensure that you are not including a mask in either dotted decimal or CIDR format. Using the “host” option implies a specific host device and therefore does not permit any mask entry.

```
ProCurve(config)# access-list 6 permit host 18.28.100.100
ProCurve(config)# access-list 6 permit host 18.28.100.100 255.255.255.255
Invalid input: 255.255.255.255
ProCurve(config)# access-list 6 permit host 18.28.100.100/32
Invalid input: 18.28.100.100/32
```

Figure C-2. Examples of Correctly and Incorrectly Specifying a Single Host

Apparent failure to log all “Deny” Matches.

Where the **log** statement is included in multiple ACEs configured with a “deny” option, a large volume of “deny” matches generating logging messages in a short period of time can impact switch performance. If it appears that the switch is not consistently logging all “deny” matches, try reducing the number of logging actions by removing the **log** statement from some ACEs configured with the “deny” action.

The switch does not allow any routed access from a specific host, group of hosts, or subnet.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL may be blocking all access by devices not specifically permitted by an entry in an ACL affecting those sources. If you are using the ACL to block specific hosts, a group of hosts, or a subnet, but want to allow any access not specifically permitted, insert **permit any** as the last explicit entry in the ACL.

The switch is not performing routing functions on a VLAN

Two possible causes of this problem are:

- Routing is not enabled. If **show running** indicates that routing is not enabled, use the **ip routing** command to enable routing.
- *On a Series 5300xl switch*, an ACL may be blocking access to the VLAN. Ensure that the switch’s IP address on the VLAN is not blocked by one of the ACE entries in an ACL applied to that VLAN. A common mistake is to either not explicitly permit the switch’s IP address as a

DA or to use a wildcard ACL mask in a deny statement that happens to include the switch's IP address. For an example of this problem, refer to the section titled "General ACL Operating Notes" in the "Access Control Lists (ACLs)" chapter of the Advanced Traffic Management Guide for your switch.

Routing Through a Gateway on the Switch Fails

Configuring a "deny" ACE that includes a gateway address can block traffic attempting to use the gateway as a next-hop.

Remote Gateway Case on a Series 5300xl Switch. For example, configuring ACL "101" (below) and applying it outbound on VLAN 1 in figure C-4 includes the router gateway (10.0.8.1) needed by devices on other networks. This can prevent the switch from sending ARP and other routing messages to the gateway router to support traffic from authorized remote networks.

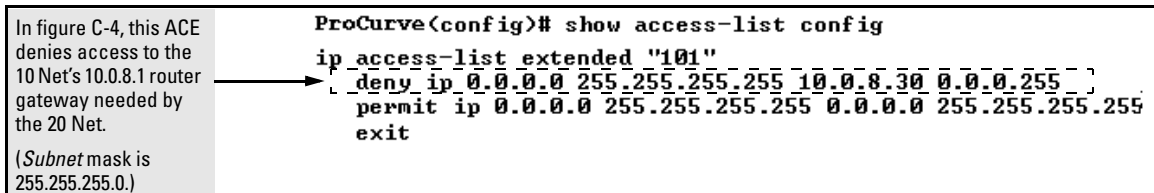


Figure C-3. Example of ACE Blocking an Entire Subnet

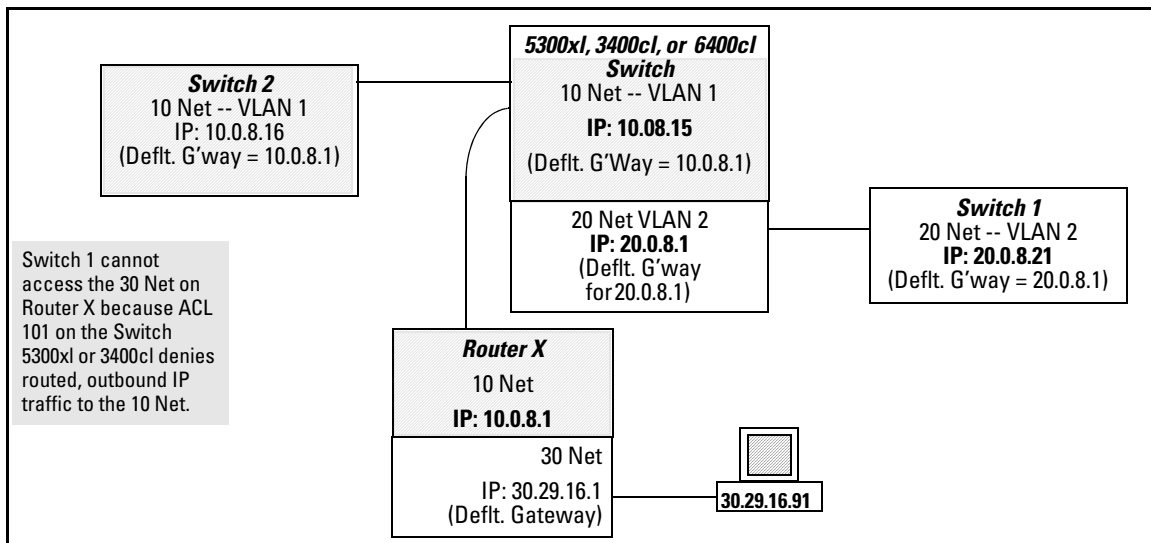


Figure C-4. Example of Inadvertently Blocking a Gateway

To avoid inadvertently blocking the remote gateway for authorized traffic from another network (such as the 20 Net in this example):

1. Configure an ACE that specifically permits authorized traffic from the remote network.
2. Configure narrowly defined ACEs to block unwanted IP traffic that would otherwise use the gateway. Such ACEs might deny traffic for a particular application, particular hosts, or an entire subnet.
3. Configure a “permit any” ACE to specifically allow any IP traffic to move through the gateway.

Local Gateway Case. If you use the switch as a gateway for traffic you want routed between subnets, use these general steps to avoid blocking the gateway for authorized applications:

1. Configure gateway security first for routing with specific permit and deny statements.
2. Permit authorized traffic.
3. Deny any unauthorized traffic that you have not already denied in step 1.

IGMP-Related Problems

IP Multicast (IGMP) Traffic That Is Directed By IGMP Does Not Reach IGMP Hosts or a Multicast Router Connected to a Port. IGMP must be enabled on the switch and the affected port must be configured for “Auto” or “Forward” operation.

IP Multicast Traffic Floods Out All Ports; IGMP Does Not Appear To Filter Traffic. The IGMP feature does not operate if the switch or VLAN does not have an IP address configured manually or obtained through DHCP/Bootp. To verify whether an IP address is configured for the switch or VLAN, do either of the following:

- **Try Using the Web Browser Interface:** If you can access the web browser interface, then an IP address is configured.
- **Try To Telnet to the Switch Console:** If you can Telnet to the switch, then an IP address is configured.
- **Using the Switch Console Interface:** From the Main Menu, check the Management Address Information screen by clicking on

1. Status and Counters

2. Switch Management Address Information

LACP-Related Problems

Unable to enable LACP on a port with the **interface < port-number > lacp** command. In this case, the switch displays the following message:

Operation is not allowed for a trunked port.

You cannot enable LACP on a port while it is configured as static **Trunk** port. To enable LACP on static-trunked port, first use the

no trunk < port-number > command to disable the static trunk assignment, then execute **interface < port-number > lacp**.

Caution

Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt your network. Before removing a port from a trunk, ProCurve recommends that you either disable the port or disconnect it from the LAN.

Mesh-Related Problems

Traffic on a dynamic VLAN does not get through the switch mesh .

GVRP enables dynamic VLANs. Ensure that all switches in the mesh have GVRP enabled. (Note that ProCurve 1600M/2400M/2424M/4000M/8000M switches do not offer GVRP. Thus, if there are any of these switches in the mesh, GVRP must be disabled for any switch in the mesh.)

The Switch Mesh Does Not Allow A ProCurve Switch 1600M/2400M/2424M/4000M/8000M Port To Join the Mesh . One of the switches in the mesh domain has detected a duplicate MAC address on multiple switches. For example:

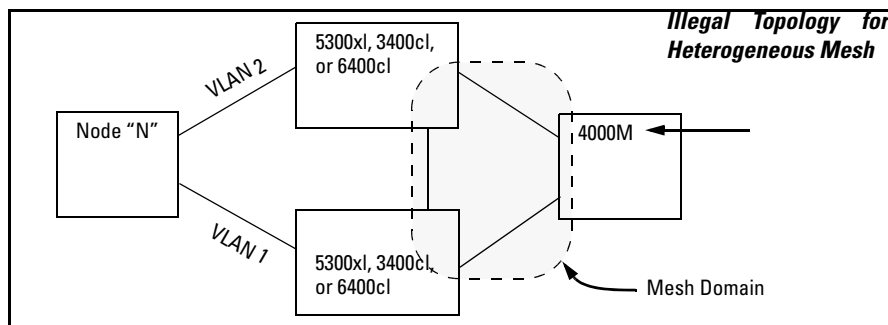


Figure C-5. Example of Illegal Topology for Heterogeneous Mesh

Changing the topology can solve this problem. Also, the duplicate MAC address must age out before the Switch 1600M/2400M/2424M/4000M/8000M port can join the mesh. Refer to the following two topics in the “Switch Meshing” chapter of the *Advanced Traffic Management Guide* for your switch:

- The section titled “Using a Heterogeneous Switch Mesh”
- The bulleted item titled “Compatibility with Older Switches” in the section titled “Requirements and Restrictions”.

Duplicate MAC Addresses on Different Switches. In a switch mesh that includes any 1600M/2400M/2424M/4000M/8000M switches, duplicate MAC addresses on different switches are not allowed. (The 1600M/2400M/2424M/4000M/8000M switches do not recognize multiple instances of a particular MAC address on different VLANs.) Refer to “The Switch Mesh Does Not Allow A ProCurve Switch 1600M/2400M/2424M/4000M/8000M Port To Join the Mesh” on page C-14.

Port-Based Access Control (802.1x)-Related Problems

Note

To list the 802.1x port-access Event Log messages stored on the switch, use **show log 802**.

See also “Radius-Related Problems” on page C-18.

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request. If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. Refer to “How 802.1x Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost. If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1x session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. Refer to “How 802.1x Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected. If **aaa authentication port-access** is configured for Local, ensure that you have entered the local *login* (operator-level) username and password of the authenticator switch into the **identity** and **secret** parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address. The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. Refer to “Note on Supplicant Statistics” in the chapter on Port-Based Access Control in the *Access Security Guide* for your switch.

The show port-access authenticator <port-list> command shows one or more ports remain open after they have been configured with control unauthorized. 802.1x is not active on the switch. After you execute **aaa port-access authenticator active**, all ports configured with **control unauthorized** should be listed as **Closed**.

```
ProCurve(config)# show port-access authenticator e A9
Port Access Authenticator Status
Port-access authenticator activated [No] : No
Access Authenticator Backend State
Port Status Control State
-----
A9 Open FU Force Auth Idle
```

PortA9 shows an "Open" status even though Access Control is set to **Unauthorized** (Force Auth). This is because the port-access authenticator has not yet been activated.

```
ProCurve(config)# aaa port-access authenticator active

ProCurve(config)# show port-access authenticator e A9
Port Access Authenticator Status
Port-access authenticator activated [No] : Yes
Access Authenticator Backend State
Port Status Control State
-----
A9 Closed FU Force Unauth Idle
```

Figure C-6. Authenticator Ports Remain "Open" Until Activated

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

```
10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key : My-Global-Key
```

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.33.18.119	1812	1813	119-only-key

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

Figure C-7. Displaying Encryption Keys

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1x configuration on that port. For example, **show port-access authenticator < port-list >** gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1x configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1x and port security either changes or is re-acquired after execution of `aaa port-access authenticator < port-list > initialize`. If the port is force-authorized with **`aaa port-access authenticator <port-list> control authorized`** command and port security is enabled on the port, then executing **`initialize`** causes the port to clear the learned address and learn a new address from the first packet it receives after you execute **`initialize`**.

A trunked port configured for 802.1x is blocked. If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-Related Problems

Loss of communication when using VLAN-tagged traffic. If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as **Untagged**.

Radius-Related Problems

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.

- Ensure that the **radius-server timeout** period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

```

10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key

  Server IP Addr  Auth Port  Acct Port  Encryption Key
  -----
  10.33.18.119    1812      1813      119-only-key
  
```

Figure C-8. Examples of Global and Unique Encryption Keys

Spanning-Tree Protocol (STP) and Fast-Uplink Problems

Caution

If you enable STP, it is recommended that you leave the remainder of the STP parameter settings at their default values until you have had an opportunity to evaluate STP performance in your network. Because incorrect STP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how STP operates. To learn the details of STP operation, refer to the IEEE 802.1D standard.

Broadcast Storms Appearing in the Network. This can occur when there are physical loops (redundant links) in the topology. Where this exists, you should enable STP on all bridging devices in the topology in order for the loop to be detected.

STP Blocks a Link in a VLAN Even Though There Are No Redundant Links in that VLAN. In 802.1Q-compliant switches STP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. Refer to “Spanning Tree Operation with VLANs” in the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.

Fast-Uplink Troubleshooting. Some of the problems that can result from incorrect usage of Fast-Uplink STP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-Uplink is configured on a switch that is the STP root device.
- Either the **Hello Time** or the **Max Age** setting (or both) is too long on one or more switches. Return the **Hello Time** and Max Age settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A “downlink” port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink STP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (Mode = **Uplink**) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup STP root switch has ports configured for fast-uplink STP and has become the root device due to a failure in the original root device.

SSH-Related Problems

Switch access refused to a client. Even though you have placed the client's public key in a text file and copied the file (using the **copy tftp pub-key-file** command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing IP SSH does not enable SSH on the switch. The switch does not have a host key. Verify by executing `show ip host-public-key`. If you see the message

```
ssh cannot be enabled until a host key is configured
(use 'crypto' command).
```

then you need to generate an SSH key pair for the switch. To do so, execute **crypto key generate**. (Refer to “2. Generating the Switch’s Public and Private Key Pair” in the SSH chapter of the *Access Security Guide* for your switch.)

Switch does not detect a client’s public key that does appear in the switch’s public key file (show ip client-public-key). The client’s public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.

```
Download failed: overlength key in key file.
Download failed: too many keys in key file.
Download failed: one or more keys is not a valid RSA
public key.
```

The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a `<CR><LF>`.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond (“hangs”) during connection phase. The switch does not support data compression in an SSH session. Clients will often have compression turned on by default, but will disable it during the negotiation phase. A client which does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned off before attempting a connection to prevent this problem.

TACACS-Related Problems

Event Log. When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All Users Are Locked Out of Access to the Switch. If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be due to how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last **write memory** command.) If you did not use **write memory** to save the authentication configuration to flash, then pressing the Reset button or cycling the power reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.
- As a last resort, use the Clear/Reset button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No Communication Between the Switch and the TACACS+ Server Application. If the switch can access the server device (that is, it can **ping** the server), then a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's **tacacs-server host** command may not be correct. (Use the switch's **show tacacs-server** command to list the TACACS+ server IP address.)

- The encryption key configured in the server does not match the encryption key configured in the switch (by using the **tacacs-server key** command). Verify the key in the server and compare it to the key configured in the switch. (Use **show tacacs-server** to list the global key. Use **show config** or **show config running** to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access Is Denied Even Though the Username/Password Pair Is Correct. Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded

For more help, refer to the documentation provided with your TACACS+ server application.

Unknown Users Allowed to Login to the Switch. Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. Refer to the documentation provided with your TACACS+ server application.

System Allows Fewer Login Attempts than Specified in the Switch Configuration. Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the **aaa authentication num-attempts** command.

TimeP, SNTP, or Gateway Problems

The Switch Cannot Find the Time Server or the Configured Gateway .

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the DEFAULT_VLAN. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-Related Problems

Monitor Port. When using the monitor port in a multiple VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized. If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link Configured for Multiple VLANs Does Not Support Traffic for One or More VLANs. One or more VLANs may not be properly configured as “Tagged” or “Untagged”. A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports. For example, VLAN_1 and VLAN_2 use the same link between switch “X” and switch “Y”.

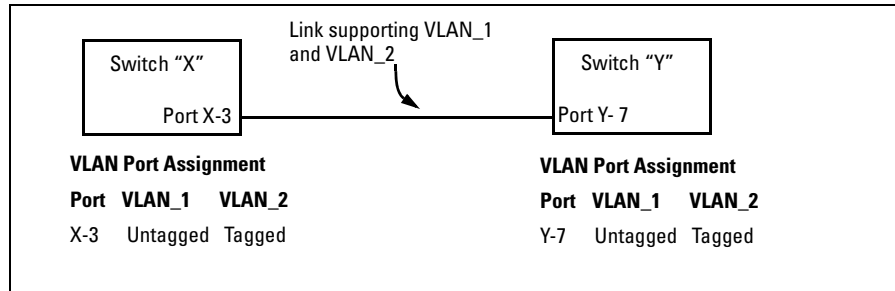


Figure C-9. Example of Correct VLAN Port Assignments on a Link

1. If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X", then it must also be configured as "Untagged" on port 7 on switch "Y". Make sure that the VLAN ID (VID) is the same on both switches.
2. Similarly, if VLAN_2 (VID=2) is configured as "Tagged on the link port on switch "A", then it must also be configured as "Tagged" on the link port on switch "B". Make sure that the VLAN ID (VID) is the same on both switches.

Duplicate MAC Addresses Across VLANs. The switches covered by this guide operate with multiple forwarding databases. Thus, duplicate MAC addresses occurring on different VLANs can appear where a device having one MAC address is a member of more than one 802.1Q VLAN, and the switch port to which the device is linked is using VLANs (instead of STP or trunking) to establish redundant links to another switch. If the other device sends traffic over multiple VLANs, its MAC address will consistently appear in multiple VLANs on the switch port to which it is linked.

Note that attempting to create redundant paths through the use of VLANs will cause problems with some switches. One symptom is that a duplicate MAC address appears in the Port Address Table of one port, and then later appears on another port. While the switches covered by this guide have multiple forwarding databases, and thus does not have this problem, some switches with a single forwarding database for all VLANs may produce the impression that a connected device is moving among ports because packets with the same MAC address but different VLANs are received on different ports. You can avoid this problem by creating redundant paths using port trunks or spanning tree.

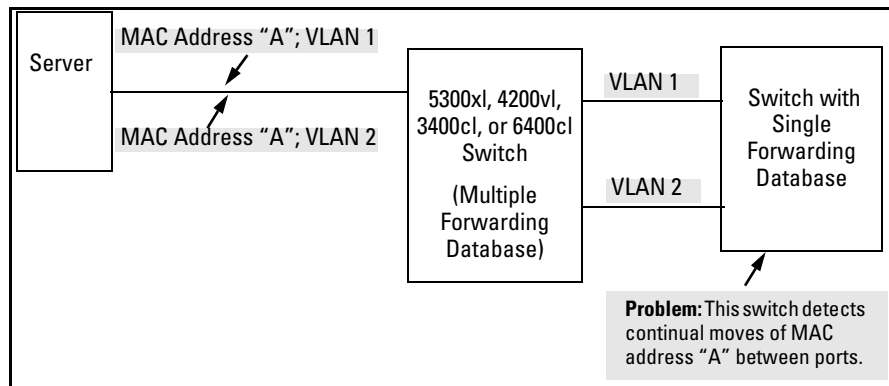


Figure C-10. Example of Duplicate MAC Address

Using the Event Log To Identify Problem Sources

The Event Log records operating events as single-line entries listed in chronological order, and serves as a tool for isolating problems. Each Event Log entry is composed of five fields:

Severity	Date	Time	System Module	Event Message
I	08/05/01	10:52:32	ports:	port A1 enabled

Severity is one of the following codes:

- I** (information) indicates routine events.
- W** (warning) indicates that a service has behaved unexpectedly.
- M** (major) indicates that a severe switch error has occurred.
- D** (debug) reserved for ProCurve internal diagnostic information.

Date is the date in *mm/dd/yy* format that the entry was placed in the log.

Time is the time in *hh:mm:ss* format that the entry was placed in the log.

System Module is the internal module (such as “ports” for port manager) that generated the log entry. If VLANs are configured, then a VLAN name also appears for an event that is specific to an individual VLAN. Table C-1 on page C-28 lists the individual modules.

Event Message is a brief description of the operating event.

The Event Log holds up to 1000 lines in chronological order, from the oldest to the newest. Each line consists of one complete event message. Once the log has received 1000 entries, it discards the current oldest line each time a new line is received. The Event Log window contains 14 log entry lines and can be positioned to any location in the log.

The Event Log will be *erased* if power to the switch is interrupted.

(The Event Log is *not* erased by using the **Reboot Switch** command in the Main Menu.)

Troubleshooting

Using the Event Log To Identify Problem Sources

Table C-1. Event Log System Modules

Module	Event Description	Module	Event Description
addrMgr	Address table	timep	Time protocol
chassis	switch hardware	udpf	UDP broadcast forwarder
bootp	bootp addressing	vlan	VLAN operations
connfilt	Connection-Rate filtering	RateLim	Rate-limiting
console	Console interface		
dhcp	DHCP addressing		
download	file transfer		
FFI	Find, Fix, and Inform -- available in the console Event Log and web browser interface alert log		
garp	GARP/GVRP		
igmp	IP Multicast		
ip	IP-related		
ipx	Novell Netware		
lACP	Dynamic LACP trunks		
ldbal	Load-Balance Protocol (meshing)		
lldp	Link-Layer Discovery Protocol		
maclock	MAC lockdown and MAC lockout		
mgr	Console management		
PIM	Protocol-Independent multicast		
ports	Change in port status; static trunks		
radius	RADIUS authentication		
snmp	SNMP communications		
ssh	Secure-Shell status		
ssl	Secure sockets layer status		
stp	Spanning Tree		
sys, system	Switch management		
telnet	Telnet activity		
tcp	Transmission control		
tftp	File transfer for new OS or config.		

Menu: Entering and Navigating in the Event Log

From the Main Menu, select **Event Log**.

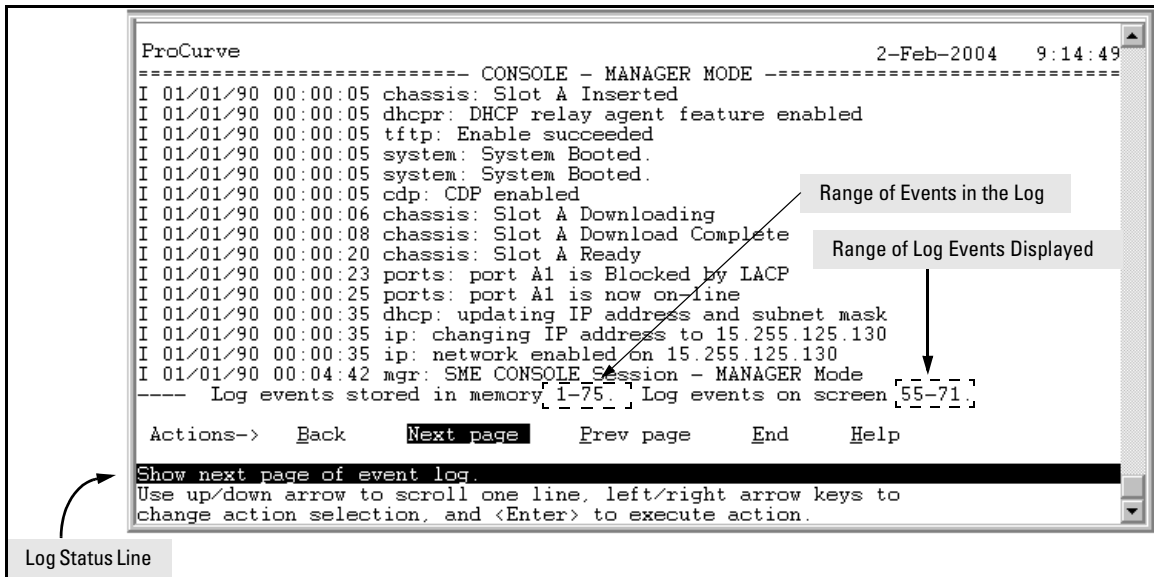


Figure C-11. Example of an Event Log Display

The *log status line* at the bottom of the display identifies where in the sequence of event messages the display is currently positioned.

To display various portions of the Event Log, either preceding or following the currently visible portion, use either the actions listed at the bottom of the display (**Next page**, **Prev page**, or **End**), or the keys described in the following table:

Table C-2. Event Log Control Keys

Key	Action
[N]	Advance the display by one page (next page).
[P]	Roll back the display by one page (previous page).
↓	Advance display by one event (down one line).
↑	Roll back display by one event (up one line).
[E]	Advance to the end of the log.
[H]	Display Help for the Event Log.

CLI: Listing Events

Syntax: show logging [-a] [<search-text>]

Uses the CLI to list:

- *Events recorded since the last boot of the switch*
- *All events recorded*
- *Event entries containing a specific keyword, either since the last boot or all events recorded*

show logging

Lists recorded log messages since last reboot.

show logging -a

Lists all recorded log messages, including those before the last reboot.

show logging -a system

Lists log messages with “system” in the text or module name.

show logging system

Lists all log messages since the last reboot that have “system” in the text or module name.

Reducing Duplicate Event Log and SNMP Trap Messages

Note

This feature is available with all software releases for the Series 3400/6400cl switches, Series 4200vl switches and with software release E.08.xx and greater on the Series 5300xl switches. Initially it applies only to Event Log messages and SNMP traps generated by the PIM software module.

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. This can flood the Event Log and any configured SNMP trap receivers with excessive, exactly identical messages. To help reduce this problem, the switch uses *log throttle periods* to regulate (throttle) duplicate messages for a given recurring event, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot. That is, when the first instance of a particular event or condition generates a message, the switch initiates a log throttle period that applies to all recurrences of that event. If the logged event recurs during the log throttle period, the switch increments the counter initiated by the first instance of the event, but does not generate a new message. If the logged event repeats again after the log throttle period expires, then the switch generates a duplicate of the first message, increments the counter, and starts a new log throttle period during which any additional instances of the event are counted, but not logged. Thus, for a particular, recurring event, the switch displays one instance of the corresponding message in the Event Log for each successive log throttle period applied to recurrences of that event. Also, each logged instance of the event message includes counter data showing how many times the event has occurred since the last reboot. The switch manages messages to SNMP trap receivers in the same way.

The log throttle period for an event depends on the event's severity level:

Severity	Log Throttle Period
I (Information)	6000 Seconds
W (Warning)	600 Seconds
M (Major)	6 Seconds
D (Debug)	60 Seconds

Example of Log Message Throttling. For example, suppose that you configure VLAN 100 on the switch to support PIM operation, but do not configure an IP address. If PIM attempted to use VLAN 100, the switch would generate the first instance of the following Event Log message and counter.

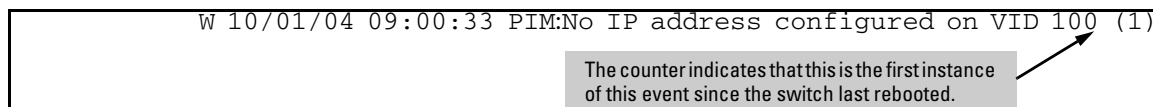


Figure C-12. Example of the First Instance of an Event Message and Counter

If PIM operation caused the same event to occur six more times during the initial log throttle period, there would be no further entries in the Event Log. However, if the event occurred again after the log throttle period expired, the switch would repeat the message (with an updated counter) and start a new log throttle period.

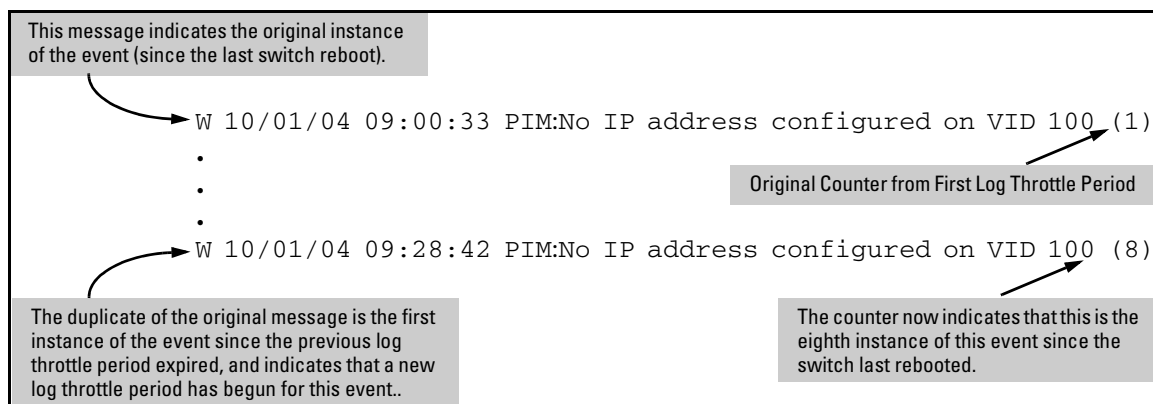


Figure C-13. Example of Duplicate Messages Over Multiple Log Throttling Periods

Note that if the same type of event occurs under different circumstances, the switch handles these as unrelated events for the purpose of Event Log messages. For example, if PIM operation simultaneously detected that VLANs 100 and 205 were configured without IP addresses, you would see log messages similar to the following:

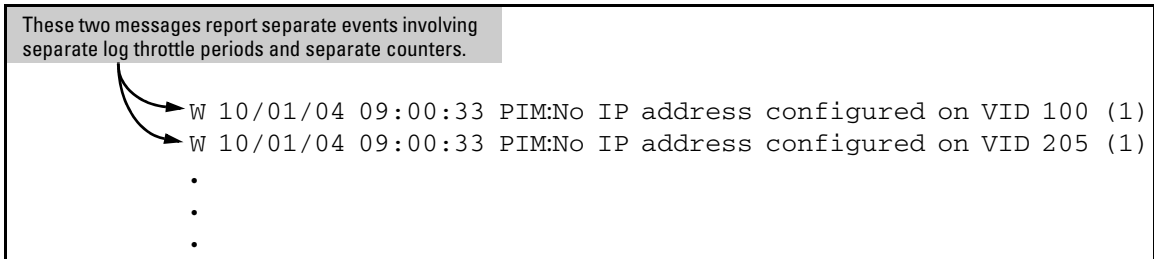


Figure C-14. Example of Log Messages Generated by Unrelated Events of the Same Type

Example of Event Counter Operation. Suppose the switch detects the following after a reboot:

- Three duplicate instances of the PIM “Send error” during the first log throttle period for this event
- Five more instances of the same Send error during the second log throttle period for this event
- Four instances of the same Send error during the third log throttle period for this event

In this case, the duplicate message would appear three times in the Event Log (once for each log throttle period for the event being described), and the Duplicate Message Counter would increment as shown in table C-3. (The same operation would apply for messages sent to any configured SNMP trap receivers.)

Table C-3. How the Duplicate Message Counter Increments

Instances During 1st Log Throttle Period	Instances During 2nd Log Throttle Period	Instances During 3rd Log Throttle Period	Duplicate Message Counter*
3			1
	5		4
		4	9

*This value always comprises the first instance of the duplicate message in the current log throttle period plus all previous occurrences of the duplicate message occurring since the switch last rebooted.

Debug and Syslog Messaging Operation

The switch’s Event Log records switch-level progress, status, and warning messages. The Debug/System-Logging (*Syslog*) feature provides a method for recording messages you can use to help in debugging network-level problems, such as routing misconfigurations and other network protocol details.

Debug enables you to specify the types of event notification messages to send to external devices. Debug messaging reports on these event types:

- ACL “deny” matches
- Selected IP routing events
- Events that generate messages for the switch’s Event Log

You can configure the switch to send debug messages to these destinations:

- Up to six Syslog servers
- A CLI session through direct RS-232 console, Telnet, or SSH

Event Notification Logging	—	Automatically sends switch-level event messages to the switch’s Event Log. Debug and Syslog do not affect this operation, but add the capability of directing Event Log messaging to an external file.
Optional Debug Commands	all	Assigns debug logging to the configured debug destination(s) for all ACL, Event Log, IP-OSPF, and IP-RIP options.
	acl	Assigns ACL Syslog logging to the debug destination(s). When there is a match with a “deny” ACE, directs the resulting message to the configured debug destination(s).
	event	Assigns standard Event Log messages to the debug destination(s). (The same messages are also sent to the switch’s Event Log, regardless of whether you enable this option.)
	IP	
	ospf	Assigns OSPF event logging to the debug destination(s).
	rip	Assigns RIP event logging to the debug destination(s).
Debug Destinations	lldp	Assigns LLDP debug logging to the debug destination(s).
	logging	Used to disable or re-enable Syslog logging if one or more Syslog servers are already configured by the separate logging < ip-addr > command. Optionally, also specifies the destination (facility) subsystem the Syslog servers must use.
	session	Assigns or re-assigns destination status to the terminal device most recently using this command to request debug output.

Figure C-15. Event Messaging Structure

Debug logging requires a logging destination (Syslog server and/or a session type), and involves the **logging** and **debug destination** commands. Actions you can perform with Debug and Syslog operation include:

Configure the switch to send Event Log messages to one or more Syslog servers. Included is the option to send the messages to the User log facility (default) on the configured server(s) or to another log facility.

Note

As of September 2004, the **logging facility** *< facility-name >* option (described on page C-40) is available on these switch models:

- Series 5300xl switches (software release E.08.xx or greater)
- Series 4200vl switches
- Series 4100gl switches (software release G.07.50 or greater)
- Series 3400cl switches
- Series 6400cl switches
- Series 2800 switches
- Series 2600 switches and the Switch 6108 (software release H.07.30 or greater)

For the latest feature information on ProCurve switches, visit the ProCurve Networking web site and check the latest release notes for the switch products you use.

-
- Configure the switch to send Event Log messages to the current management-access session (serial-connect CLI, Telnet CLI, or SSH).
 - Disable all Syslog debug logging while retaining the Syslog addresses from the switch configuration. This allows you to configure Syslog messaging and then disable and re-enable it as needed.
 - Display the current debug configuration. If Syslog logging is currently active, this includes the Syslog server list.
 - Display the current Syslog server list when Syslog logging is disabled.

Debug Command Operation

As shown in figure C-15, the **debug** command performs two main functions:

- Specifies the type(s) of event messaging to send to a destination.
- Specifies the destination(s) of the selected message types.

Except as noted below, rebooting the switch returns the debug destination and debug message types to their default settings (disabled).

Note

Using the **logging < dest-ip-addr >** command to configure a Syslog server address creates an exception to the above general operation. Refer to “Syslog Operation” on page C-39.

Debug Types

This section describes the types of debug messages the switch can send to configured debug destinations.

Syntax: [no] debug < debug-type >

acl

*When a match occurs on an ACL “deny” Access Control Entry (with **log** configured), the switch sends an ACL message to the configured debug destination(s). For more on ACLs, refer to the chapter titled “Access Control Lists” in the Advanced Traffic Management Guide for your switch. (Default: Disabled)*

all

Configures the switch to send all debug types to the configured debug destination(s). (Default: Disabled)

event

*Configures the switch to send Event Log messages to the configured debug destination(s). **Note:** This has no effect on event notification messages the switch routinely sends to the Event Log itself. Also, this debug type is automatically enabled in these cases:*

- *If there is currently no Syslog server address configured and you use **logging < ip-addr >** to configure an address.*
- *If there is currently at least one Syslog server address configured and the switch is rebooted or reset.*

ip

Enables all IP-OSPF message types for the configured destinations.

lldp

Enables all LLDP message types for the configured destinations.

Syntax: [no] debug < debug-type > (Continued)

ip [ospf < adj | event | flood | lsa-generation | packet | retransmission
| spf >]

For the configured debug destination(s):

ospf < adj | event | flood | lsa-generation | packet | retransmission
| spf > — *Enables the specified IP-OSPF message type.*

adj — *Adjacency changes.*

event — *OSPF events.*

flood — *Information on flood messages.*

lsa-generation — *New LSAs added to database.*

packet — *Packets sent/received.*

retransmission — *Retransmission timer messages.*

spf — *Path recalculation messages.*

ip [rip < database | event | trigger >]

rip < database | event | trigger > — *Enables the specified RIP
message type for the configured destination(s).*

database — *Display database changes.*

event — *Display RIP events.*

trigger — *Display trigger messages.*

(Default: Event (log) message type.)

Debug Destinations

Debug enables you to disable and re-enable Syslog messaging to configured servers, and to designate a CLI session to receive messaging of any debug type.

Syntax: [no] debug destination < logging | session >

logging

*This command enables Syslog logging to the configured Syslog server(s). That is, the switch sends the debug message types (specified by the **debug < debug-type >** command in the preceding section) to the configured Syslog server(s). (Default: Logging disabled)*

(To configure a Syslog server IP address, refer to “Syslog Operation” on page C-39.)

Note: Debug messages from Series 5300xl switches running software release E.07.21 or greater, Series 4200vl switches, and any 3400cl/6400cl switches carry a “debug” severity level. Because some Syslog servers, in their default configuration, ignore Syslog messages with this severity level, you should ensure that the Syslog servers you intend to receive debug messages are configured to accept the “debug” severity level. For more information, refer to “Operating Notes for Debug and Syslog” on page C-44.

session

*Enables or disables transmission of event notification messages to the CLI session that most recently executed this command. The session can be on any one terminal emulation device with serial, Telnet, or SSH access to the CLI at the Manager level prompt (**ProCurve#_**). If more than one terminal device has a console session with the CLI, you can redirect the destination from the current device to another device. Do so by executing **debug destination session** in the CLI on the terminal device on which you now want to display event messages.*

*Event message types received on the selected CLI session are those specified by the **debug < debug-type >** command. (Refer to “Debug Types” on page C-36.)*

Syslog Operation

Syslog is a client-server logging tool that allows a client switch to send event notification messages to a networked device operating with Syslog server software. Messages sent to a Syslog server can be stored to a file for later debugging analysis. Use of Syslog requires that you set up a Syslog server application on a networked host accessible to the switch. (Refer to the documentation for the Syslog server application you select.) Except as described below, you must use the **debug** command to specify the message types the switch sends to the configured Syslog server(s).

Syntax: [no] logging < syslog-ip-addr >

Enables or disables Syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (Syslog) and the Event debug type. Thus, at a minimum, the switch begins sending Event Log messages to the configured Syslog server(s). The ACL, IP-OSPF, and/or IP-RIP message types will also be sent to the Syslog server(s) if they are currently enabled debug types. (Refer to “Debug Types” on page C-36.)

no logging *removes all currently configured Syslog logging destinations from the switch.*

no logging < syslog-ip-address > *removes only the specified Syslog logging destination from the switch.*

If you use the “no” form of the command to delete the only remaining logging address, debug destination logging is disabled on the switch, but the Event debug type is not changed from its current setting.

*To block messages to the configured Syslog server(s) from any currently enabled debug type, use **no debug < debug-type >**. (Refer to “Debug Types” on page C-36.)*

*To disable Syslog logging on the switch without deleting the server addresses, use **no debug destination logging**. Note that, unlike the case where there are no Syslog servers configured, if one or more Syslog servers are already configured, but Syslog messaging is disabled, adding a new server address to those already configured does not re-enable Syslog messaging. In this case, you must use **debug destination logging** to re-enable Syslog messaging.*

Syntax: [no] logging facility < facility-name >

The logging facility specifies the destination subsystem the Syslog server(s) must use. (All configured Syslog servers must use the same subsystem.) ProCurve recommends the default (user) subsystem unless your application specifically requires another subsystem. Options include:

user (the default) — Random user-level messages
kern — Kernel messages
mail — Mail system
daemon — System daemons
auth — Security/Authorization messages
syslog — Messages generated internally by Syslog
lpr — Line-Printer subsystem
news — Netnews subsystem
uucp — uucp subsystem
cron — cron/at subsystem
sys9 — cron/at subsystem
sys10 - sys14 — Reserved for system use
local10 - local17 — Reserved for system use

For a listing of applicable ProCurve switches, refer to the Note on page C-35.

Viewing the Debug Configuration

Syntax: show debug

This command displays the currently configured debug logging destination(s) and type(s). For examples of show debug output, refer to figure C-16 on page C-42.

Steps for Configuring Debug and Syslog Messaging

1. Skip this step if you do not want to use a Syslog server.

If you want a Syslog server as a destination for debug messaging:

- a. Use this command to configure the Syslog server IP address and enable Syslog logging:

```
ProCurve(config)# logging < ip-addr >
```

Using this command when there are no Syslog server IP addresses already configured enables both debug messaging to a Syslog server and the Event debug-type, which means that the switch begins sending Event Log messages to the server, regardless of other debug types that may be configured.

- b. Use the command in step “a” to configure any additional Syslog servers you want to use, up to a total of six. (When multiple server IP addresses are configured, the switch sends the selected debug message types to all such addresses.)
- c. If you want Event Log messages sent to the Syslog server, skip this step. Otherwise, use this command to block Event Log messages to the server:

```
ProCurve# no debug event
```

2. If you do not want a CLI session for a destination, skip this step.

Otherwise, from the device to which you want the switch to send debug messages:

- a. Use a serial, Telnet, or SSH connection to access the switch's CLI.
- b. Execute this command:

```
ProCurve# debug destination session
```

3. Enable the debug types for which you want messages sent to the Syslog server(s) and/or the current session device:

```
ProCurve# debug < acl | all | event | ip [ospf-opt]>
```

Repeat this step if necessary to enable multiple debug types.

Example: Suppose that there are no Syslog servers configured on the switch (the default). Configuring one Syslog server enables debug logging to that server and also enables Event Log messages to be sent to the server.

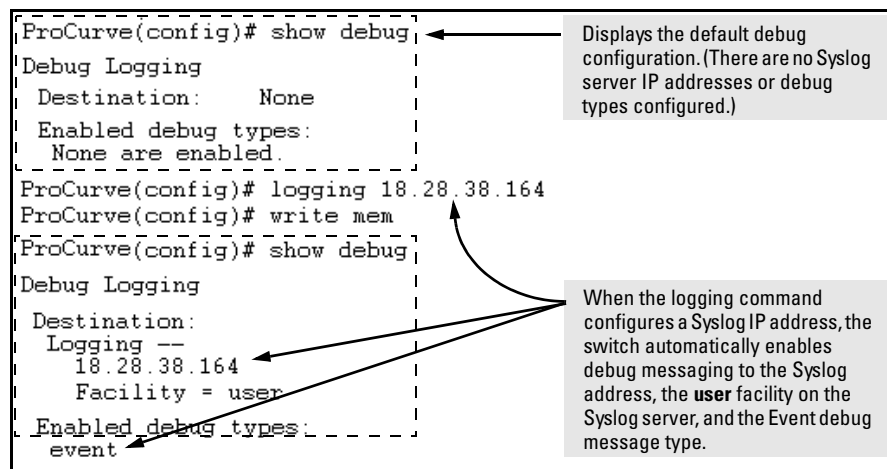


Figure C-16. Example of Configuring Basic Syslog Operation

Note that after you enable Syslog logging, if you do not want Event Log messages sent to the Syslog server(s), you can block such messages by executing **no debug event**. (This has no effect on standard logging of messages in the switch's Event Log.)

Example. Suppose that you want to:

- Configure Syslog logging of ACL and IP-OSPF packet messages on a Syslog server at 18.38.64.164 (with **user** as the default logging facility).
- Also display these messages in the CLI session of your terminal device's management access to the switch.
- Prevent the Switch's standard Event Log messages from going to the Syslog server and the CLI.

Assuming the debug/Syslog feature is disabled on the switch, you would use the commands shown in figure C-17 to configure the above operation.

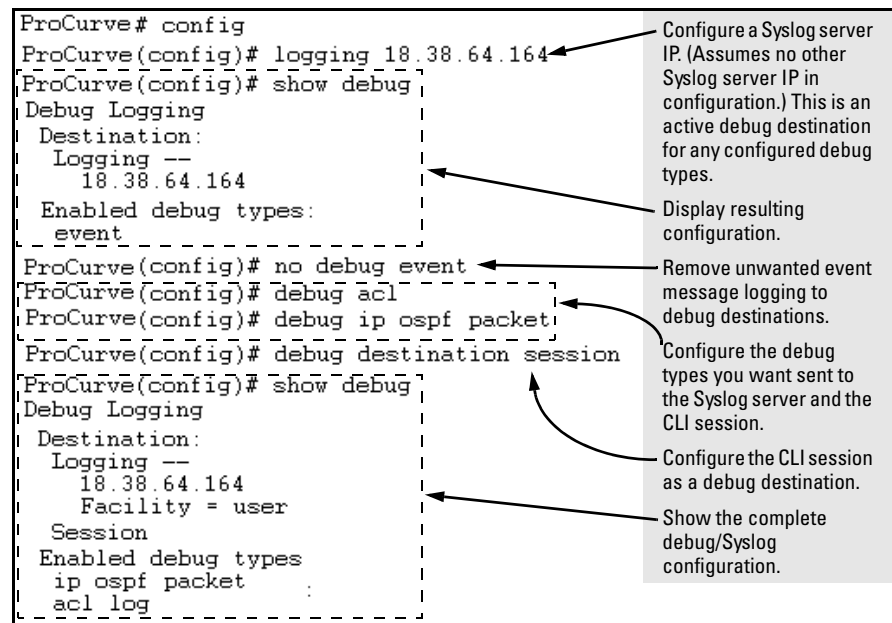


Figure C-17. Example Debug/Syslog Configuration for Multiple Types and Destinations

Operating Notes for Debug and Syslog

- **Rebooting the Switch or pressing the Reset button resets the Debug Configuration.**

Debug Option	Effect of a Reboot or Reset
logging (destination)	If any Syslog server IP addresses are in the startup-config file, they are saved across a reboot and the logging destination option remains enabled. Otherwise, the logging destination is disabled.
Session (destination)	Disabled.
ACL (event type)	Disabled.
All (event type)	Disabled.
Event (event type)	If a Syslog server is configured in the startup-config file, resets to enabled, regardless of prior setting. Disabled if no Syslog server is configured.
IP (event type)	Disabled.

- **Debug commands do not affect message output to the Event Log.**
As a separate option, invoking debug with the **event** option causes the switch to send Event Log messages to whatever debug destination(s) you configure (session and/or logging), as well as to the Event Log.
- **Ensure that your Syslog server(s) will accept Debug messages.** All Syslog messages resulting from debug operation carry a “debug” severity. If you configure the switch to transmit debug messages to a Syslog server, ensure that the server’s Syslog application is configured to accept the “debug” severity level. (The default configuration for some Syslog applications ignores the “debug” severity level.)

Diagnostic Tools

Diagnostic Features

Feature	Default	Menu	CLI	Web
Port Auto negotiation	n/a	n/a	n/a	n/a
Ping Test	n/a	—	page C-48	page C-47
Link Test	n/a	—	page C-48	page C-47
Display Config File	n/a	—	page C-50	page C-50
Admin. and Troubleshooting Commands	n/a	—	page C-52	—
Factory-Default Config	page C-57 (Buttons)	—	page C-57	—
Port Status	n/a	pagesB-9and B-10	pagesB-9and B-10	pagesB-9and B-10

Port Auto-Negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to **Auto** mode.
2. If the attached end-node does not have an **Auto** mode setting, then you must manually configure the switch port to the same setting as the end-node port. Refer to Chapter 10, “Port Status and Basic Configuration”.

Ping and Link Tests

The Ping test and the Link test are point-to-point tests between your switch and another IEEE 802.3-compliant device on your network. These tests can tell you whether the switch is communicating properly with another device.

Note

To respond to a Ping test or a Link test, the device you are trying to reach must be IEEE 802.3-compliant.

Ping Test. This is a test of the path between the switch and another device on the same or another IP network that can respond to IP packets (ICMP Echo Requests).

Link Test. This is a test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.

Web: Executing Ping or Link Tests

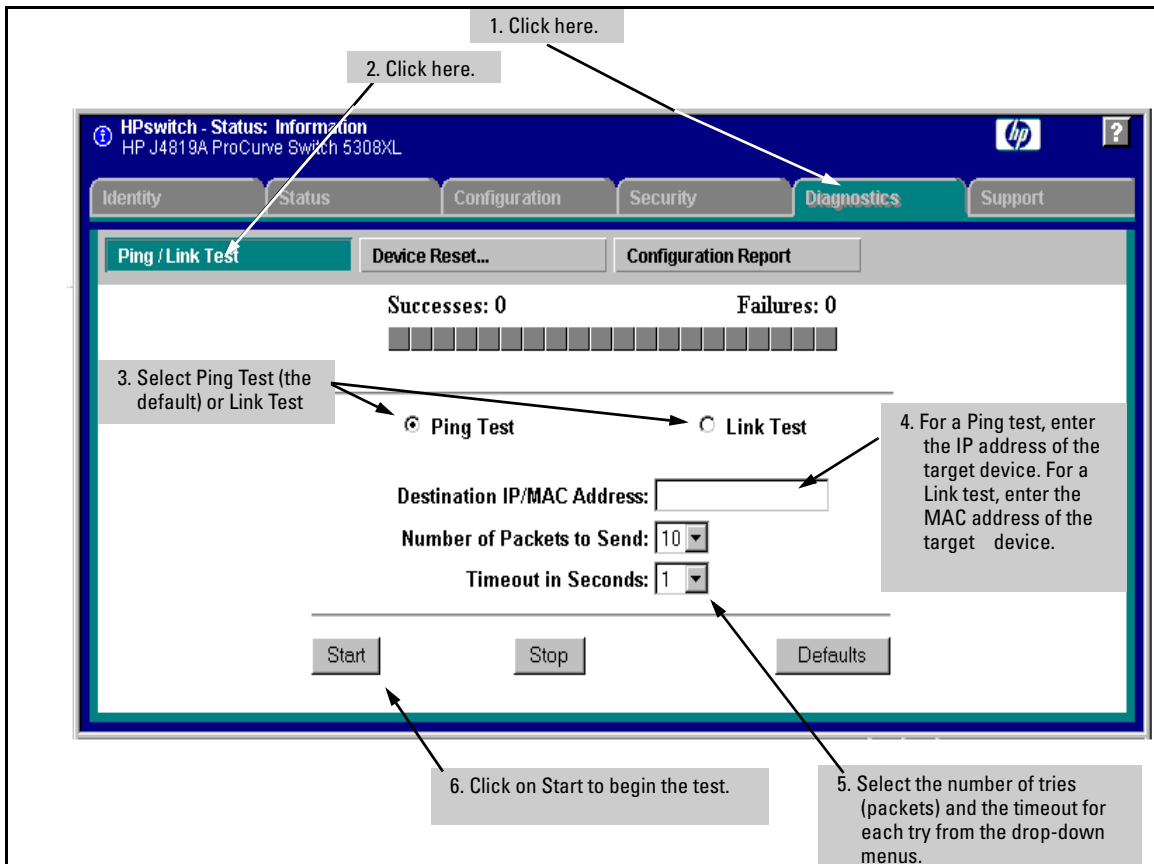


Figure C-18. Link and Ping Test Screen on the Web Browser Interface

Successes indicates the number of Ping or Link packets that successfully completed the most recent test.

Failures indicates the number of Ping or Link packets that were unsuccessful in the last test. Failures indicate connectivity or network performance problems (such as overloaded links or devices).

Destination IP/MAC Address is the network address of the target, or destination, device to which you want to test a connection with the switch. An IP address is in the X.X.X.X format where X is a decimal number between 0 and 255. A MAC address is made up of 12 hexadecimal digits, for example, 0060b0-080400.

Number of Packets to Send is the number of times you want the switch to attempt to test a connection.

Timeout in Seconds is the number of seconds to allow per attempt to test a connection before determining that the current attempt has failed.

To halt a Link or Ping test before it concludes, click on the Stop button.
To reset the screen to its default settings, click on the Defaults button.

CLI: Ping or Link Tests

Ping Tests. You can issue single or multiple ping tests with varying repetitions and timeout periods. The defaults and ranges are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: ping < ip-address > [repetitions < 1 - 999 >] [timeout < 1 - 256 >]

Basic Ping Operation	→	ProCurve> ping 10.28.227.103 10.28.227.103 is alive, time = 15 ms
Ping with Repetitions	→	ProCurve> ping 10.28.227.103 repetitions 3 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 15 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping with Repetitions and Timeout	→	ProCurve> ping 10.28.227.103 repetitions 3 timeout 2 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 10 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping Failure	→	ProCurve> ping 10.28.227.105 Target did not respond.

Figure C-19. Examples of Ping Tests

To halt a ping test before it concludes, press **[Ctrl] [C]**.

Link Tests. You can issue single or multiple link tests with varying repetitions and timeout periods. The defaults are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: link < mac-address > [repetitions < 1 - 999 >] [timeout < 1 - 256 >]
[vlan < vlan-id >]

Basic Link Test	HP4108#link 0030c1-7fcc40 Link-test passed.
Link Test with Repetitions	ProCurve# link 0030c1-7fcc40 repetitions 3 802.2 TEST packets sent: 3, responses received: 3
Link Test with Repetitions and Timeout	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN; Test Fail	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 222 802.2 TEST packets sent: 3, responses received: 0

Figure C-20. Example of Link Tests

Displaying the Configuration File

The complete switch configuration is contained in a file that you can browse from either the web browser interface or the CLI. It may be useful in some troubleshooting scenarios to view the switch configuration.

CLI: Viewing the Configuration File

Using the CLI, you can display either the running configuration or the startup configuration. (For more on these topics, see appendix C, “Switch Memory and Configuration”.)

Syntax: write terminal

Displays the running configuration.

show config

Displays the startup configuration.

show running-config

Displays the running-config file.

Web: Viewing the Configuration File

To display the running configuration, through the web browser interface:

1. Click on the **Diagnostics** tab.
2. Click on **[Configuration Report]**
3. Use the right-side scroll bar to scroll through the configuration listing.

Listing Switch Configuration and Operation Details

The **show tech** command outputs, in a single listing, switch operating and running configuration details from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration
- Event Log listing
- Boot History
- Port settings
- Status and counters — port status

- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

Syntax: show tech

Executing **show tech** outputs a data listing to your terminal emulator. However, using your terminal emulator's text capture features, you can also save **show tech** data to a text file for viewing, printing, or sending to an associate. For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the show tech output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

To Copy show tech output to a Text File. This example uses the Microsoft Windows terminal emulator. To use another terminal emulator application, refer to the documentation provided with that application.

1. In Hyperterminal, click on **Transfer | Capture Text...**

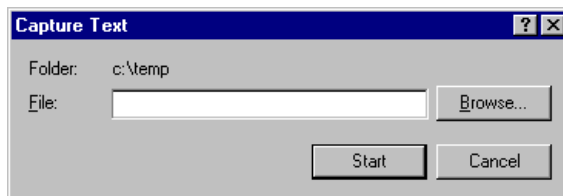


Figure C-21. The Capture Text window of the Hyperterminal Application

2. In the **File** field, enter the path and file name under which you want to store the **show tech** output.

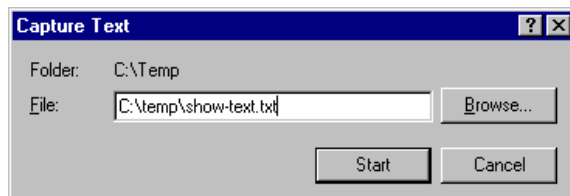


Figure C-22. Example of a Path and Filename for Creating a Text File from show tech Output

3. Click **[Start]** to create and open the text file.

4. Execute **show tech**:

```
ProCurve# show tech
```

- a. Each time the resulting listing halts and displays -- MORE --, press the Space bar to resume the listing.
- b. When the CLI prompt appears, the show tech listing is complete. At this point, click on **Transfer | Capture Text | Stop** in HyperTerminal to stop copying data into the text file created in the preceding steps.

Note

Remember to do the above step to stop HyperTerminal from copying into the text file. Otherwise, the text file remains open to receiving additional data from the HyperTerminal screen.

5. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

CLI Administrative and Troubleshooting Commands

These commands provide information or perform actions that you may find helpful in troubleshooting operating problems with the switch.

Note

For more on the CLI, refer to chapter 3, “Using the Command Line Reference (CLI)”.

Syntax: show version

Shows the software version currently running on the switch, and the flash image from which the switch booted (primary or secondary).

show boot-history

Displays the switch shutdown history.

show history

Displays the current command history.

Syntax: show version

[no] page

Toggles the paging mode for display commands between continuous listing and per-page listing.

setup

Displays the Switch Setup screen from the menu interface.

repeat

Repeatedly executes the previous command until a key is pressed.

kill

Terminates all other active sessions.

Traceroute Command

The **traceroute** command enables you to trace the route from the switch to a host address.

This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute **traceroute**, it uses the same default settings unless you specify otherwise for that instance of the command.

Syntax: traceroute < ip-address >

Lists the IP address of each hop in the route, plus the time in microseconds for the **traceroute** packet reply to the switch for each hop.

To halt an ongoing traceroute search, press the **[Ctrl] [C]** keys.

[minttl < 1-255 >]

*For the current instance of **traceroute**, changes the minimum number of hops allowed for each probe packet sent along the route. If **minttl** is greater than the actual number of hops, then the output includes only the hops at and above the **minttl** threshold. (The hops below the threshold are not listed.) If **minttl** matches the actual number of hops, only that hop is shown in the output. If **minttl** is less than the actual number of hops, then all hops are listed. For any instance of **traceroute**, if you want a **minttl** value other than the default, you must specify that value. (Default: 1)*

[maxttl < 1-255 >]

*For the current instance of **traceroute**, changes the maximum number of hops allowed for each probe packet sent along the route. If the destination address is further from the switch than **maxttl** allows, then **traceroute** lists the IP addresses for all hops it detects up to the **maxttl** limit. For any instance of **traceroute**, if you want a **maxttl** value other than the default, you must specify that value. (Default: 30)*

[timeout < 1-120 >]

*For the current instance of **traceroute**, changes the timeout period the switch waits for each probe of a hop in the route. For any instance of **traceroute**, if you want a **timeout** value other than the default, you must specify that value. (Default: 5 seconds)*

[probes < 1-5 >]

*For the current instance of **traceroute**, changes the number of queries the switch sends for each hop in the route. For any instance of **traceroute**, if you want a **probes** value other than the default, you must specify that value. (Default: 3)*

A Low Maxttl Causes Traceroute To Halt Before Reaching the Destination Address. For example, executing **traceroute** with its default values for a destination IP address that is four hops away produces a result similar to this:

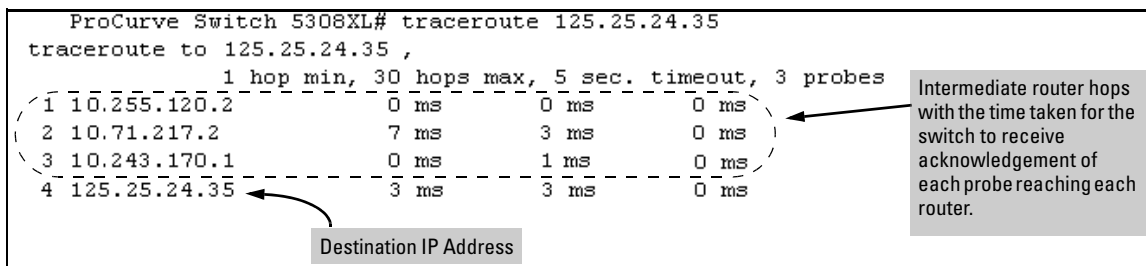


Figure C-23. Example of a Completed Traceroute Enquiry

Continuing from the previous example (figure C-23, above), executing **traceroute** with an insufficient **maxttl** for the actual hop count produces an output similar to this:

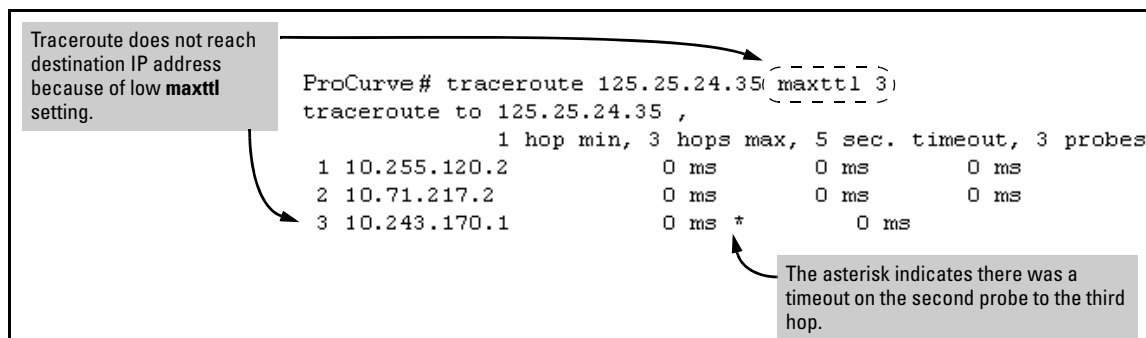


Figure C-24. Example of Incomplete Traceroute Due to Low Maxttl Setting

If A Network Condition Prevents Traceroute from Reaching the Destination. Common reasons for Traceroute failing to reach a destination include:

- Timeouts (indicated by one asterisk per probe, per hop; see figure C-24, above.)
- Unreachable hosts
- Unreachable networks
- Interference from firewalls
- Hosts configured to avoid responding

Executing traceroute where the route becomes blocked or otherwise fails results in an output marked by timeouts for all probes beyond the last detected hop. For example with a maximum hop count of 7 (**maxttl** = 7), where the route becomes blocked or otherwise fails, the output appears similar to this:

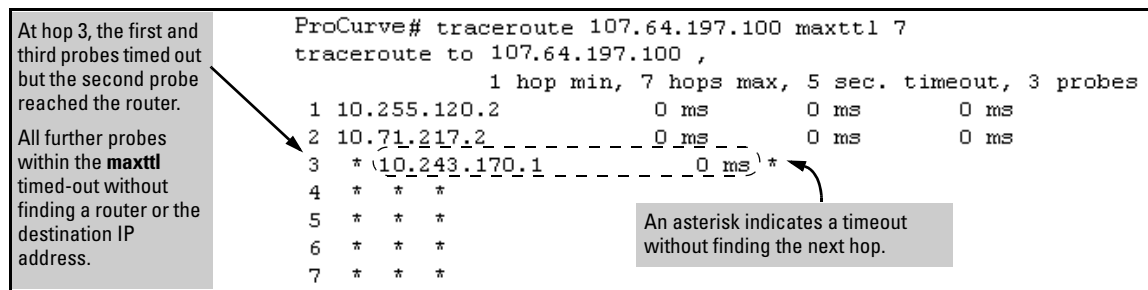


Figure C-25. Example of Traceroute Failing to Reach the Destination Address

Restoring the Factory-Default Configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process momentarily interrupts the switch operation, clears any passwords, clears the console Event Log, resets the network counters to zero, performs a complete self test, and reboots the switch into its factory default configuration including deleting an IP address. There are two methods for resetting to the factory-default configuration:

- CLI
- Clear/Reset button combination

Note

ProCurve recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem, to a directly connected PC.

CLI: Resetting to the Factory-Default Configuration

This command operates at any level *except* the Operator level.

Syntax: erase startup-configuration

Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.

Note

The **erase startup-config** command does not clear passwords.

Clear/Reset: Resetting to the Factory-Default Configuration

To execute the factory default reset, perform these steps:

1. Using pointed objects, simultaneously press both the Reset and Clear buttons on the front of the switch.

2. Continue to press the Clear button while releasing the Reset button.
3. When the Self Test LED begins to flash, release the Clear button.

The switch will then complete its self test and begin operating with the configuration restored to the factory default settings.

Restoring a Flash Image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the **erase flash** command to erase a good OS image file from the opposite flash location.

To Recover from an Empty or Corrupted Flash State. Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch.

Note

The following procedure requires the use of Xmodem, and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

1. Start the terminal emulator program.
2. Ensure that the terminal program is configured as follows:
 - Baud rate: 9600
 - 1 stop bit
 - No parity
 - No flow control
 - 8 Bits

3. Use the Reset button to reset the switch. The following prompt should then appear in the terminal emulator:

```
Enter h or ? for help.
```

```
=>
```

4. Since the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For example:

- a. Change the switch baud rate to 115,200 Bps.

```
=> sp 115200
```

- b. Change the terminal emulator baud rate to match the switch speed:
 - i. In HyperTerminal, select **Call | Disconnect**.
 - ii. Select **File | Properties**.
 - iii. click on **[Configure .]**.
 - iv. Change the baud rate to **115200**.
 - v. Click on **[OK]**. In the next window, click on **[OK]** again.
 - vi. Select **Call | Connect**
 - vii. Press **[Enter]** one or more times to display the => prompt.

5. Start the Console Download utility by typing **do** at the => prompt and pressing **[Enter]**:

```
=> do
```

6. You will then see this prompt:

```
You have invoked the console download utility.  
Do you wish to continue? (Y/N)>_
```

7. At the above prompt:
 - a. Type **y** (for Yes)
 - b. Select **Transfer | File** in HyperTerminal.
 - c. Enter the appropriate filename and path for the OS image.
 - d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).
 - e. Click on **[Send]**.

If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:

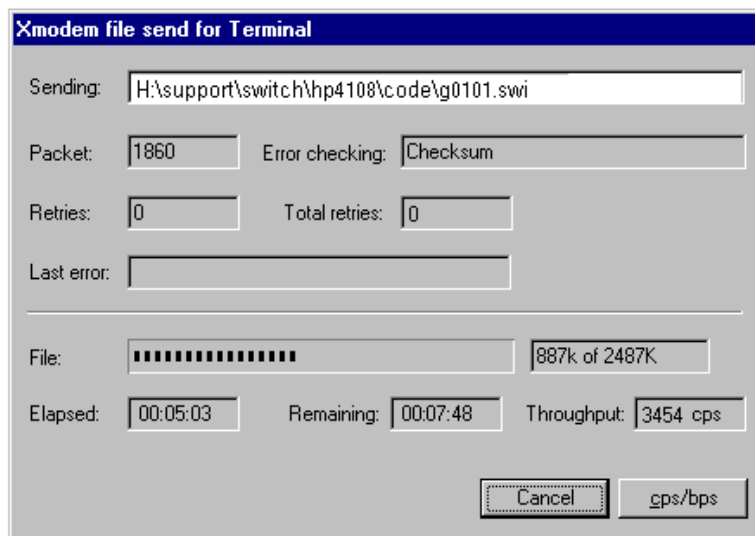


Figure C-26. Example of Xmodem Download in Progress

8. When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.

MAC Address Management

Contents

Overview	D-2
Determining MAC Addresses	D-3
Menu: Viewing the Switch's MAC Addresses	D-4
CLI: Viewing the Port and VLAN MAC Addresses	D-5
Viewing the MAC Addresses of Connected Devices	D-8

Overview

The switch assigns MAC addresses in these areas:

- For management functions, one Base MAC address is assigned to the default VLAN (VID = 1). (All VLANs on the switches covered in this guide use the same MAC address.)
- For internal switch operations: One MAC address per port (See “CLI: Viewing the Port and VLAN MAC Addresses” on page D-5.)

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.

Note

The switch’s base MAC address is also printed on a label affixed to the switch.

Determining MAC Addresses

MAC Address Viewing Methods

Feature	Default	Menu	CLI	Web
view switch's base (default vlan) MAC address and the addressing for any added VLANs	n/a	D-4	D-5	—
view port MAC addresses (hexadecimal format)	n/a	—	D-5	—

- **Use the menu interface** to view the switch's base MAC address and the MAC address assigned to any VLAN you have configured on the switch. (The same MAC address is assigned to VLAN1 and all other VLANs configured on the switch.)

Note

The switch's base MAC address is used for the default VLAN (VID = 1) that is always available on the switch. This is true for dynamic VLANs as well; the base MAC address is the same across all VLANs.

- **Use the CLI** to view the switch's port MAC addresses in hexadecimal format.

Menu: Viewing the Switch's MAC Addresses

The Management Address Information screen lists the MAC addresses for:

- Base switch (default VLAN; VID = 1)
- Any additional VLANs configured on the switch.

Also, the Base MAC address appears on a label on the back of the switch.

Note

The Base MAC address is used by the first (default) VLAN in the switch. This is usually the VLAN named “DEFAULT_VLAN” unless the name has been changed (by using the VLAN Names screen). On the switches covered by this guide, the VID (VLAN identification number) for the default VLAN is always “1”, *and cannot be changed*.

To View the MAC Address (and IP Address) assignments for VLANs Configured on the Switch:

1. From the Main Menu, Select

1. Status and Counters

2. Switch Management Address Information

If the switch has only the default VLAN, the following screen appears. If the switch has multiple static VLANs, each is listed with its address data.

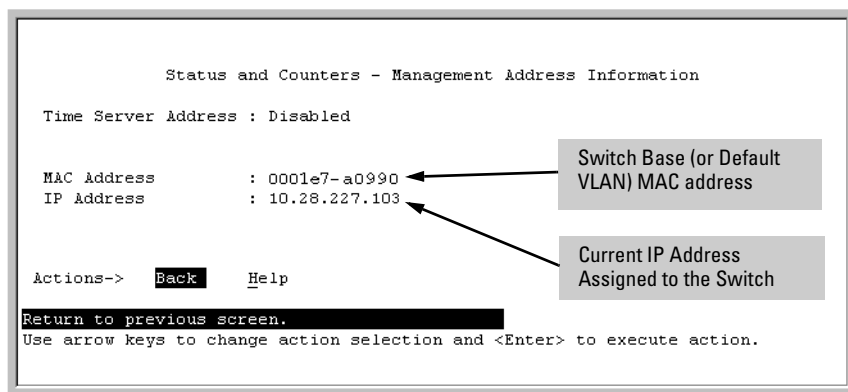


Figure D-1. Example of the Management Address Information Screen

CLI: Viewing the Port and VLAN MAC Addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the spanning-tree protocol. Using the **walkmib** command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.

Switch Series	MAC Address Allocation
5300xl and 4200vl	The switch allots 26 MAC addresses per slot. For a given slot, if a four-port module is installed, then the switch uses the first four MAC addresses in the allotment for that slot, and the remaining 22 MAC addresses are unused. If a 24-port module is installed, the switch uses the first 24 MAC addresses in the allotment, and so-on. The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the walkmib listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)
3400cl-24G	This switch uses 24 MAC addresses for the 24 fixed ports, plus two MAC addresses if an optional 10-gigabit expansion module is installed. The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the walkmib listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)
3400cl-48G	Same as for the 3400cl-24G, except this model uses 48 MAC addresses for the 48 fixed ports.
6400cl	These switches use 6 MAC addresses for the 6 fixed ports, plus two MAC addresses if an optional 10-gigabit expansion module is installed. The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the walkmib listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)

To display the switch's MAC addresses, use the **walkmib** command at the command prompt:

Note

This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

1. If the switch is at the CLI Operator level, use the **enable** command to enter the Manager level of the CLI.
2. Type the following command to display the MAC address for each port on the switch:

```
ProCurve# walkmib ifPhysAddress
```

(The above command is not case-sensitive.)

For example, with a 4-port module in slot A of a 5304xl switch, a 24-port module in slot B, and four nondefault VLANs configured:

<pre> ProCurve(config)# walkmib_ifphysaddress ifPhysAddress.1 = 00 04 ea 5e 20 ff ifPhysAddress.2 = 00 04 ea 5e 20 fe ifPhysAddress.3 = 00 04 ea 5e 20 fd ifPhysAddress.4 = 00 04 ea 5e 20 fc ifPhysAddress.27 = 00 04 ea 5e 20 e5 ifPhysAddress.28 = 00 04 ea 5e 20 e4 ifPhysAddress.29 = 00 04 ea 5e 20 e3 ifPhysAddress.30 = 00 04 ea 5e 20 e2 ifPhysAddress.31 = 00 04 ea 5e 20 e1 ifPhysAddress.32 = 00 04 ea 5e 20 e0 ifPhysAddress.33 = 00 04 ea 5e 20 df ifPhysAddress.34 = 00 04 ea 5e 20 de ifPhysAddress.35 = 00 04 ea 5e 20 dd ifPhysAddress.36 = 00 04 ea 5e 20 dc ifPhysAddress.37 = 00 04 ea 5e 20 db ifPhysAddress.38 = 00 04 ea 5e 20 da ifPhysAddress.39 = 00 04 ea 5e 20 d9 ifPhysAddress.40 = 00 04 ea 5e 20 d8 ifPhysAddress.41 = 00 04 ea 5e 20 d7 ifPhysAddress.42 = 00 04 ea 5e 20 d6 ifPhysAddress.43 = 00 04 ea 5e 20 d5 ifPhysAddress.44 = 00 04 ea 5e 20 d4 ifPhysAddress.45 = 00 04 ea 5e 20 d3 ifPhysAddress.46 = 00 04 ea 5e 20 d2 ifPhysAddress.47 = 00 04 ea 5e 20 d1 ifPhysAddress.48 = 00 04 ea 5e 20 d0 ifPhysAddress.49 = 00 04 ea 5e 20 cf ifPhysAddress.50 = 00 04 ea 5e 20 ce ifPhysAddress.282 = 00 04 ea 5e 20 00 ifPhysAddress.381 = 00 04 ea 5e 20 00 ifPhysAddress.431 = 00 04 ea 5e 20 00 ifPhysAddress.456 = 00 04 ea 5e 20 00 ifPhysAddress.481 = 00 04 ea 5e 20 00 ifPhysAddress.4377 = ProCurve(config)# </pre>	<div> ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A (Addresses 5 - 24 in slot A are unused, and addresses 25 and 26 are reserved.) </div> <div> ifPhysAddress.27 - 50: Ports B1 - B24 in Slot B (Addresses 51 - 52 in slot B are reserved.) </div> <div> ifPhysAddress.282 Base MAC Address (MAC Address for default VLAN; VID = 1) </div> <div> ifPhysAddress.381, 431, 456, and 481 Physical addresses for non-default VLANs configured on the switch. On the switches covered by this guide, all VLANs use the same MAC address as the Default VLAN. Refer to "Multiple VLAN Considerations" in the "Static Virtual LANs (VLANs)" chapter of the <i>Advanced Traffic Management Guide</i> for your switch. </div>
---	---

Figure D-2. Example of Port MAC Address Assignments on a Series 5300xl Switch

ProCurve Switch 2824(config)# walkmib ifphysaddress	
ifPhysAddress.1 = 00 08 83 08 db 3f	
ifPhysAddress.2 = 00 08 83 08 db 3e	
ifPhysAddress.3 = 00 08 83 08 db 3d	
ifPhysAddress.4 = 00 08 83 08 db 3c	
ifPhysAddress.5 = 00 08 83 08 db 3b	
ifPhysAddress.6 = 00 08 83 08 db 3a	
ifPhysAddress.7 = 00 08 83 08 db 39	
ifPhysAddress.8 = 00 08 83 08 db 38	
ifPhysAddress.9 = 00 08 83 08 db 37	
ifPhysAddress.10 = 00 08 83 08 db 36	ifPhysAddress.1 - 24: Ports 1 - 24 (A 3400cl-48G switch includes addresses 1 - 48 for the fixed ports.)
ifPhysAddress.11 = 00 08 83 08 db 35	
ifPhysAddress.12 = 00 08 83 08 db 34	
ifPhysAddress.13 = 00 08 83 08 db 33	
ifPhysAddress.14 = 00 08 83 08 db 32	
ifPhysAddress.15 = 00 08 83 08 db 31	
ifPhysAddress.16 = 00 08 83 08 db 30	
ifPhysAddress.17 = 00 08 83 08 db 2f	
ifPhysAddress.18 = 00 08 83 08 db 2e	
ifPhysAddress.19 = 00 08 83 08 db 2d	
ifPhysAddress.20 = 00 08 83 08 db 2c	
ifPhysAddress.21 = 00 08 83 08 db 2b	ifPhysAddress.25 - 26: Ports 25 - 26 (Addresses 25 - 26 appear only if a 10-gigabit expansion module is installed in the switch. On a 3400cl-48G switch, these ports are numbered 49 and 50.)
ifPhysAddress.22 = 00 08 83 08 db 2a	
ifPhysAddress.23 = 00 08 83 08 db 29	
ifPhysAddress.24 = 00 08 83 08 db 28	
ifPhysAddress.25 = 00 08 83 08 db 27	ifPhysAddress.102 Base MAC Address (MAC Address for default VLAN; VID = 1)
ifPhysAddress.26 = 00 08 83 08 db 26	
ifPhysAddress.102 = 00 08 83 08 db 20	
ifPhysAddress.4197 =	

Figure D-3. Example of Port MAC Address Assignments on a 3400cl-24G Switch

Viewing the MAC Addresses of Connected Devices

Syntax: show mac-address [| mac-addr |

Lists the MAC addresses of the devices the switch has detected, along with the number of the specific port on which each MAC address was detected.

[port-list]

Lists the MAC addresses of the devices the switch has detected, on the specified port(s).

[mac-addr]

Lists the port on which the switch detects the specified MAC address. Returns the following message if the specified MAC address is not detected on any port in the switch:

MAC address <mac-addr> not found.

[vlan < vid >]

Lists the MAC addresses of the devices the switch has detected on ports belonging to the specified VLAN, along with the number of the specific port on which each MAC address was detected.

To list the MAC addresses of devices the switch has detected, use the **show mac-address** command.

Daylight Savings Time on ProCurve Switches

This information applies to the following ProCurve switches:

- 212M
- 224M
- 1600M
- 2400M
- 2424M
- 4000M
- 8000M
- Series 2500
- Series 2600
- Series 2800
- Series 3400cl
- Series 4100gl
- Series 4200vl
- Series 5300xl
- Series 6400cl
- 6108
- AdvanceStack Switches
- AdvanceStack Routers

ProCurve switches provide a way to automatically adjust the system clock for Daylight Savings Time (DST) changes. To use this feature you define the month and date to begin and to end the change from standard time. In addition to the value “none” (no time changes), there are five pre-defined settings, named:

- Alaska
- Canada and Continental US
- Middle Europe and Portugal
- Southern Hemisphere
- Western Europe

The pre-defined settings follow these rules:

Alaska:

- Begin DST at 2am the first Sunday on or after April 24th.
- End DST at 2am the first Sunday on or after October 25th.

Canada and Continental US:

- Begin DST at 2am the first Sunday on or after April 1st.
- End DST at 2am the first Sunday on or after October 25th.

Middle Europe and Portugal:

- Begin DST at 2am the first Sunday on or after March 25th.
- End DST at 2am the first Sunday on or after September 24th.

Southern Hemisphere:

- Begin DST at 2am the first Sunday on or after October 25th.
- End DST at 2am the first Sunday on or after March 1st.

Western Europe:

- Begin DST at 2am the first Sunday on or after March 23rd.
- End DST at 2am the first Sunday on or after October 23rd.

A sixth option named “User defined” allows you to customize the DST configuration by entering the beginning month and date plus the ending month and date for the time change. The menu interface screen looks like this (all month/date entries are at their default values):

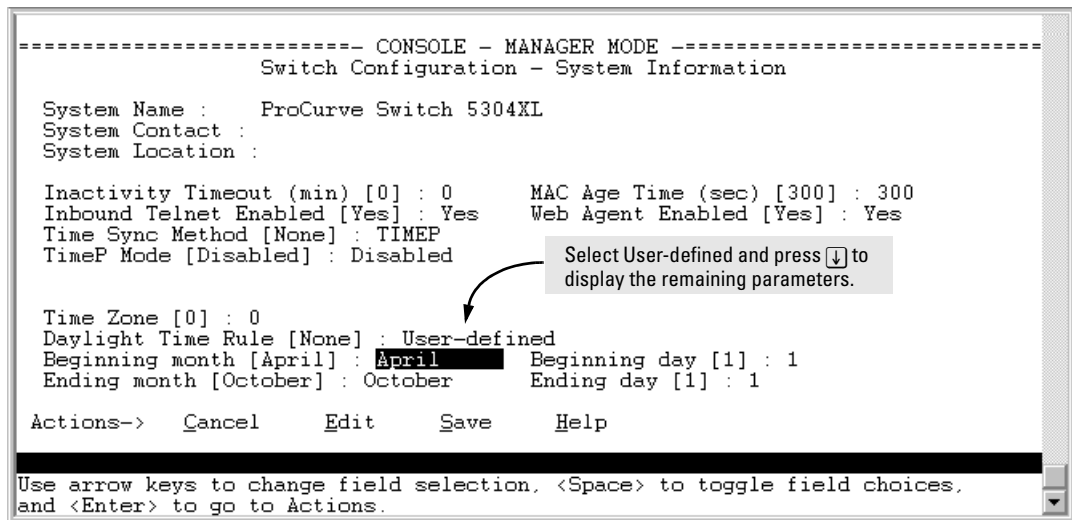


Figure E-1. Menu Interface with “User-Defined” Daylight Time Rule Option

Before configuring a “User defined” Daylight Time Rule, it is important to understand how the switch treats the entries. The switch knows which dates are Sundays, and uses an algorithm to determine on which date to change the system clock, given the configured “Beginning day” and “Ending day”:

- If the configured day is a Sunday, the time changes at 2am on that day.
- If the configured day is not a Sunday, the time changes at 2am on the first Sunday after the configured day.

This is true for both the “Beginning day” and the “Ending day”.

With that algorithm, one should use the value “1” to represent “first Sunday of the month”, and a value equal to “number of days in the month minus 6” to represent “last Sunday of the month”. This allows a single configuration for every year, no matter what date is the appropriate Sunday to change the clock.

— This page is intentionally unused. —

Index

Symbols

=> prompt ... C-58

Numerics

802.1x

LLDP blocked ... 15-32

802.1X effect, LLDP ... 15-71

A

access

manager ... 15-12

operator ... 15-12

Access Contrtoller Module

5300xl features not supported ... 12-8

access control server ... 12-4

Access Controller Module

BIOS POST event log messages ... 12-32

configuring client VLANs ... 12-19

configuring on network ... 12-14

configuring uplink ports ... 12-19

displaying status ... 12-24

downloading software ... 12-30

features of ... 12-7

general operation ... 12-3

managing ... 12-27

module operation ... 12-5

network address translation ... 12-11

operating rules ... 12-14

overview ... 12-5

resetting to factory defaults ... 12-30

routing support ... 12-10

uplink ports ... 12-6

VLAN base ... 12-18

VLANs ... 12-12

ACL

debug ... C-34, C-36

See also debug command.

gateway fails ... C-11

troubleshooting ... C-8

ACL resources ... 14-5

ACM, *see* Access Controller Module

Actions line ... 3-9, 3-10, 3-11

location on screen ... 3-9

address table, port ... B-13

address, network manager ... 15-4, 15-5

advertise location ... 15-48

alert log ... 5-20

alert types ... 5-21

disabling ... 5-25

setting the sensitivity level ... 5-24

sorting the entries ... 5-20

arp age, default ... 8-7

asterisk ... 3-10, 3-13, 6-27

asterisk, in traceroute ... C-55

authentication trap ... 15-19, 15-22

See also SNMP.

authentication trap, configuring ... 15-22

authorized IP managers

SNMP, blocking ... 15-3

auto MDI/MDI-X configuration, display ... 10-16

auto MDI/MDI-X operation ... 10-15, 10-16

auto MDI/MDI-X port mode, display ... 10-16

auto negotiation ... 10-4

Auto-10 ... 13-4, 13-7, 13-18

autonegotiate ... 15-48

B

bandwidth

displaying utilization ... 5-17

BIOS POST messages ... 12-32

boot

See also reboot.

boot command ... 6-4, 6-18

boot ROM console ... A-4

boot ROM mode ... C-58

boot, from primary flash ... 6-19

Bootp ... 8-2, 8-12

Bootp table file ... 8-14

Bootptab file ... 8-13

effect of no reply ... C-7

operation ... 8-13

See also DHCP.

using with Unix systems ... 8-13

Bootp/DHCP differences ... 8-13

Bootp/DHCP, LLDP ... 15-44

broadcast limit ... 10-5, 10-14

broadcast storm ... 13-3, C-20

broadcast traffic, IPX ... 10-5, 10-14

browser interface

See web browser interface.

C

CDP

configuration, viewing ... 15-75

data collection ... 15-73

default CDP operation ... 15-73

disabled ... 15-73

general operation ... 15-74

LLDP neighbor data ... 15-72

mappings to LLDP data fields ... 15-73

neighbor devices ... 15-73

on individual ports ... 15-77

read-only operation ... 15-72, 15-74

Clear + Reset button combination ... 6-35

Clear button ... 5-10

restoring factory default configuration ... C-58

CLI

context level ... 10-10

moving to or from the menu ... 4-7

client traffic ... 12-11

client VLANs ... 12-12

command line interface

See CLI.

communities, SNMP ... 15-13

viewing and configuring with the CLI ... 15-14

viewing and configuring with the menu ... 15-13

conceptual view ... 12-5

configuration ... 3-7

Bootp ... 8-13

comparing startup to running ... 6-6

console ... 7-3

copying ... A-23

download ... A-4

factory default ... 6-9, 8-2

IP ... 8-2

network monitoring ... B-23

permanent ... 6-7

permanent change defined ... 6-5

port ... 10-1

port trunk group ... 13-1

port, duplex ... 10-9

port, speed ... 10-9

quick ... 3-8

reboot to activate ... 3-13

restoring factory defaults ... C-57

saving from menu interface ... 3-10

serial link ... 7-3

SNMP ... 15-4, 15-5, 15-11

SNMP communities ... 15-13, 15-14

startup ... 3-10

system ... 7-9

Telnet access configuration ... 7-3

transferring ... A-23

trap receivers ... 15-19

viewing ... 6-6

web browser access ... 7-3

configuration file

browsing for troubleshooting ... C-50

multiple

configuration file, multiple

after first E.09.xx reboot ... 6-26

applicable software release ... 6-22, 6-23

applicable switch models ... 6-22

applications ... 6-23

asterisk ... 6-27

backupConfig ... 6-24

change policy ... 6-29

Clear + Reset button combination ... 6-35

copy from tftp host ... 6-37

copy to tftp host ... 6-37

create new file ... 6-25, 6-32, 6-33

current file in use ... 6-27

default reboot from primary ... 6-30

effect of E.09.xx download ... 6-28

erasing ... 6-34

factory default, E.09.xx installed ... 6-28

memory assignments ... 6-26

memory slot ... 6-24, 6-27, 6-30

minconfig ... 6-30, 6-34

newconfig ... 6-30

oldConfig ... 6-26, 6-28

override reboot policy ... 6-29

policy, override ... 6-31

power cycle ... 6-30

pre E.09.xx in one flash ... 6-26

primary boot path ... 6-27

reboot policy options ... 6-24

reboot policy, override ... 6-29

reboot process ... 6-25

reload ... 6-31

rename config file ... 6-32

reset ... 6-30

- running-config file ... 6-25
- running-config file operation ... 6-24
- secondary boot path ... 6-27
- sftp/scp transfer ... 6-39
- show config file content ... 6-29
- show multiple files ... 6-27
- slot 1, use ... 6-39
- startup-config ... 6-24
- startup-config file ... 6-25
- transition to multiple files ... 6-26
- Unable to copy ... 6-32
- workingConfig ... 6-24, 6-26, 6-28
- xmodem from host ... 6-38
- xmodem to host ... 6-38
- configuring uplink VLAN ... 12-19
- console ... C-7
 - configuring ... 7-3
 - ending a session ... 3-5
 - features ... 2-3
 - Main menu ... 3-7
 - navigation ... 3-9, 3-10
 - operation ... 3-10
 - starting a session ... 3-4
 - status and counters access ... 3-7
 - troubleshooting access problems ... C-5
- context level
 - global config ... 8-10
- copy
 - multiple config file, tftp ... 6-37
- CPU utilization ... B-6

D

- date format ... C-27
- date, configure ... 7-13
- debug
 - overview ... C-34
 - session options ... C-34
- debug command
 - "debug" severity and Syslog servers ... C-38, C-44
 - configuring messaging ... C-40
 - destinations ... C-34, C-38
 - event ... C-36
 - event log ... C-34, C-44
 - operating notes ... C-44
 - OSPF ... C-37
 - RIP ... C-37
 - session configuration ... C-38

- show debug ... C-40
- structure ... C-34
- syntax ... C-36
- debug logging, LLDP ... 15-30
- default gateway ... 8-3
- default trunk type ... 13-10
- default VLAN ... 12-12
- Device Passwords Window ... 5-8
- DHCP ... 8-12
 - address problems ... C-7
 - effect of no reply ... C-7
 - manual gateway precedence ... 8-12
- DHCP/Bootp differences ... 8-13
- DHCP/Bootp process ... 8-12
- DHCP/Bootp, LLDP ... 15-44
- diagnostics tools ... C-45
 - browsing the configuration file ... C-50
 - ping and link tests ... C-45
- display commands
 - show commands ... 12-24
- displaying duplex information ... 15-64
- DNS name ... 5-4
- Domain Name Server ... 5-4
- downlink ports ... 12-6
- download ... A-4
 - See also* switch software.
 - switch-to-switch ... A-18
 - troubleshooting ... A-21
 - Xmodem ... A-16
- download software ... A-18
- download, TFTP ... A-4, A-5
- duplex advertisements ... 15-46
- duplicate MAC address
 - See* MAC address.
- Dyn1
 - See* LACP.

E

- edge ports ... 14-4
- Emergency Location Id Number ... 15-57
- Emergency Location Identification Number ... 15-25
- ending a console session ... 3-5
- event log ... 3-7, C-27
 - navigation ... C-29
 - severity level ... C-27
 - use during troubleshooting ... C-27
 - with debug ... C-34, C-44

excessive packets ... 14-34

F

factory default configuration

restoring ... 6-9, C-57

failure, switch software download ... A-22

fault detection ... 5-8

fault detection policy ... 5-8, 5-24

fault detection policy, setting ... 5-24

fault detection window ... 5-24

fault-tolerance ... 13-5

filter, source-port ... 14-33

firmware version ... B-6

flash memory ... 3-10, 6-3

flow control

constraints ... 10-5, 10-11

global ... 10-10, 10-11

global requirement ... 10-5, 10-10

jumbo packets ... 14-28, 14-32

per-port ... 10-5, 10-10, 10-11

flow control, effect on rate-limiting ... 14-7, 14-18

flow control, status ... B-10

flow control, terminal ... 7-3

format, date ... C-27

format, time ... C-27

friendly port names

See port names, friendly.

G

gateway ... 8-3, 8-5, 8-12

routing fails ... C-11

gateway (IP) address ... 8-4, 8-6

gateway, manual config priority ... 8-12

gateway, on primary VLAN ... 8-4

giant packets ... 14-34

global config level ... 8-10

GMB

See guaranteed minimum bandwidth.

guaranteed minimum bandwidth

apportioning unallocated bandwidth ... 14-23

configuration ... 14-23

described ... 14-21

displaying configuration ... 14-6

displaying current configuration ... 14-25

granularity of bandwidth settings ... 14-26

operating notes ... 14-26

operation ... 14-21

outbound queue priority ... 14-22

starving queues ... 14-23

H

Help ... 3-11, 5-13

Help line, about ... 3-9

Help line, location on screens ... 3-9

help, online inoperable ... 5-13

hop, router ... 8-10

HP Auto-MDIX feature ... 10-15

I

IEEE 802.1AB/D9 ... 15-26

IEEE 802.1d ... C-19

IEEE 802.3ab ... 10-4

IEEE P802.1AB/D9 ... 15-31

IGMP

host not receiving ... C-13

not working ... C-13

statistics ... B-19

inactivity timeout ... 7-4

Inbound Telnet Enabled parameter ... C-6

Inconsistent value ... 15-40

installing ... 12-6

invalid input ... 4-12

IP ... 8-7

CLI access ... 8-6

configuration ... 8-2

DHCP/Bootp ... 8-2

duplicate address ... C-7

duplicate address, DHCP network ... C-7

effect when address not used ... 8-11

features available with and without ... 8-11

gateway ... 8-3

gateway (IP) address ... 8-4

menu access ... 8-5

multiple addresses in VLAN ... 8-3, 8-8

subnet ... 8-3, 8-8

subnet mask ... 8-2, 8-6

time server address ... 9-10, 9-19

Time-To-Live ... 8-7, 8-10

TTL ... 8-7, 8-10

using for web browser interface ... 5-4

web access ... 8-10

IP address

- for SNMP management ... 15-3
- multiple in a VLAN ... 8-8
- removing or replacing ... 8-10
- IP preserve
 - DHCP server ... 8-16
 - overview ... 8-15
 - rules, operating ... 8-16
 - summary of effect ... 8-18
- IPX
 - network number ... B-7
- IPX broadcast traffic ... 10-5, 10-14

J

- Java ... 5-4, 5-5
- jumbo packets
 - configuration ... 14-29
 - excessive inbound ... 14-32
 - flow control ... 14-28, 14-32
 - GVRP operation ... 14-28
 - management VLAN ... 14-32
 - maximum size ... 14-27
 - meshing ... 14-28
 - MTU ... 14-27
 - port adds and moves ... 14-28
 - port speed ... 14-28
 - security concerns ... 14-33
 - standard MTU ... 14-27
 - switch mesh domain ... 14-34
 - through non-jumbo ports ... 14-33
 - traffic sources ... 14-28
 - troubleshooting ... 14-34
 - VLAN tag ... 14-27
 - voice VLAN ... 14-32

K

- key components ... 12-5
- kill command ... 7-8

L

- LACP
 - 802.1x, not allowed ... 13-21
 - active ... 13-15
 - CLI access ... 13-11
 - default port operation ... 13-20
 - described ... 13-6, 13-18
 - Dyn1 ... 13-7

- dynamic ... 13-19
- enabling dynamic trunk ... 13-15
- full-duplex required ... 13-4, 13-18
- IGMP ... 13-23
- monitoring static trunk ... B-23
- no half-duplex ... 13-23
- operation not allowed ... C-13
- overview of port mode settings ... 13-5
- passive ... 13-15
- removing port from active trunk ... 13-16
- restrictions ... 13-21
- standby link ... 13-19
- status, terms ... 13-21
- STP ... 13-23
- trunk limit ... 13-19
- VLANs ... 13-23
- with 802.1x ... 13-21
- with port security ... 13-22
- learning bridge ... 8-2
- limit, broadcast ... 10-14
- link speed, port trunk ... 13-3
- link test ... C-46
 - for troubleshooting ... C-45
- link, serial ... 7-3
- LLDP
 - 802.1D-compliant switch ... 15-71
 - 802.1x blocking ... 15-32
 - 802.1X effect ... 15-71
 - active port ... 15-25
 - adjacent device ... 15-25
 - advertisement ... 15-25
 - advertisement content ... 15-43
 - advertisement data ... 15-62
 - advertisement, mandatory data ... 15-43
 - advertisement, optional data ... 15-44
 - advertisements, delay interval ... 15-39
 - CDP neighbor data ... 15-72
 - chassis ID ... 15-43
 - chassis type ... 15-43
 - clear statistics counters ... 15-68
 - comparison with CDP data fields ... 15-73
 - configuration options ... 15-28
 - configuring optional data ... 15-44
 - data options ... 15-29
 - data read options ... 15-30
 - data unit ... 15-26
 - debug logging ... 15-30
 - default ... 15-73

- default configuration ... 15-33
- DHCP/Bootp operation ... 15-32
- disable, per-port ... 15-42
- display neighbor data ... 15-65
- ELIN ... 15-25
- enable/disable, global ... 15-37
- features ... 15-24
- general operation ... 15-27
- global counters ... 15-68
- holdtime multiplier ... 15-39
- hub, packet-forwarding ... 15-28
- IEEE 802.1AB/D9 ... 15-26
- IEEE P802.1AB/D9 ... 15-31
- Inconsistent value ... 15-40
- information options ... 15-29
- invalid frames ... 15-69
- IP address advertisement ... 15-31
- IP address subelement ... 15-43
- IP address, advertisement ... 15-71
- IP address, DHCP/Bootp ... 15-44
- IP address, options ... 15-43
- IP address, version advertised ... 15-43
- Link-Layer Discovery Protocol ... 15-24
- LLDP-aware ... 15-26
- LLDPDU ... 15-26
- mandatory TLVs ... 15-71
- MIB ... 15-27, 15-31
- neighbor ... 15-26
- neighbor data remaining ... 15-71
- neighbor data, displaying ... 15-65
- neighbor statistics ... 15-68
- neighbor, maximum ... 15-70
- on 3400/6400cl switches ... 15-32
- operating rules ... 15-31
- operation ... 15-27
- optional data, configuring ... 15-45
- outbound packet options ... 15-29
- packet boundaries ... 15-27
- packet dropped ... 15-28
- packet time-to-live ... 15-30
- packet-forwarding ... 15-27, 15-71
- packets not forwarded ... 15-26
- per-port counters ... 15-69
- port description ... 15-44
- port ID ... 15-43
- port speed ... 15-46
- port trunks ... 15-31
- port type ... 15-43
- refresh interval ... 15-38
- reinitialization delay ... 15-40
- remote management address ... 15-30
- remote manager address ... 15-43
- reset counters ... 15-68
- rxonly ... 15-42
- setmib, delay interval ... 15-39
- setmib, reinit delay ... 15-41
- show advertisement data ... 15-62
- show commands ... 15-34, 15-36
- show outbound advertisement ... 15-63
- SNMP notification ... 15-29
- SNMP traps ... 15-29
- spanning-tree blocking ... 15-32
- standards compatibility ... 15-31
- statistics ... 15-68
- statistics, displaying ... 15-68
- system capabilities ... 15-44
- system description ... 15-44
- system name ... 15-44
- terminology ... 15-25
- time-to-live ... 15-28, 15-38
- TLV ... 15-27
- transmission frequency ... 15-28
- transmission interval, change ... 15-38
- transmit and receive ... 15-28
- transmit/receive modes ... 15-28
- transmit/receive modes, per-port ... 15-42
- trap notice interval ... 15-42
- trap notification ... 15-41
- trap receiver, data change notice ... 15-41
- TTL ... 15-28, 15-30
- txonly ... 15-42
- VLAN, untagged ... 15-71
- walkmib ... 15-31
- LLDPDU ... 15-26
- LLDP-MED
 - displaying speed ... 15-64
 - ELIN ... 15-57
 - enable or disable ... 15-28
 - endpoint support ... 15-48
 - fast start control ... 15-51
 - location data ... 15-56
 - medTlvenable ... 15-53
 - Neighbors MIB ... 15-65
 - topology change notification ... 15-50
 - Voice over IP ... 15-47
- load balancing

See port trunk.
logging, command ... C-36
logical port ... 13-8
loop, network ... 13-3
lost password ... 5-10

M

MAC address ... 8-13, B-6, D-2
 displaying detected devices ... D-8
 duplicate ... C-20, C-25
 learned ... B-13
 per slot ... D-5
 per-slot or per-port ... D-5
 port ... D-2, D-4
 same MAC, multiple VLANs ... D-6
 switch ... D-2
 VLAN ... D-2, D-5
 walkmib ... D-5
management
 interfaces described ... 2-2
 server URL ... 5-12, 5-13
 server URL default ... 5-13
management VLAN
 See VLAN.
manager access ... 15-12
manager password ... 5-8, 5-10
Manual, IP address ... 8-6
MDI/MDI-X configuration, display ... 10-16
MDI/MDI-X operation ... 10-15
MDI/MDI-X port mode, display ... 10-16
media type, port trunk ... 13-3
memory
 flash ... 3-10, 6-3
 startup configuration ... 3-10
menu interface
 configuration changes, saving ... 3-10
 moving to or from the CLI ... 4-7
mesh
 jumbo packets ... 14-34
 monitoring ... B-23
meshed ports, monitoring ... B-23
MIB ... 15-4
MIB listing ... 15-4
MIB, HP proprietary ... 15-4
MIB, standard ... 15-4
mirroring
 See port monitoring.

MLTS ... 15-26
monitoring
 See port monitoring.
monitoring meshed ports ... B-23
monitoring traffic ... B-23
monitoring, port ... B-23
Multiline Telephone system ... 15-26
multinetting ... 8-3, 8-8
 See also ACLs.
multinetting, limit ... 8-8
multiple configuration file
 See configuration file, multiple.
multiple forwarding database ... B-7, B-16
multiple VLAN ... 15-3
multi-port bridge ... 8-2

N

NANP ... 15-26
navigation, console interface ... 3-9, 3-10
navigation, event log ... C-29
network management functions ... 15-5, 15-13
network manager address ... 15-4, 15-5
network monitoring
 traffic overload ... B-23
Network Monitoring Port screen ... B-23
network slow ... C-7
North American Numbering Plan ... 15-26

O

online help ... 5-13
online help location ... 5-13
operating system
 See switch software.
operation not allowed, LACP ... C-13
operator access ... 15-12
operator password ... 5-8, 5-10
OS
 See switch software.
 version ... A-19
OSPF
 debug command ... C-37
out-of-band ... 2-4

P

password ... 5-8, 5-10
 creating ... 5-8

- delete ... 3-7, 5-10
- if you lose the password ... 5-10
- lost ... 5-10
- manager ... 5-8
- operator ... 5-8
- set ... 3-7
- setting ... 5-9
- using to access browser and console ... 5-10
- PCM/PCM+
 - starting web browser ... 5-4
- PD ... 15-26
- ping test ... C-46
 - for troubleshooting ... C-45
- PoE ... 15-47
 - 802.3af ... 11-6
 - 802.3af standard ... 11-4
 - advertisements ... 15-56
 - cable ... 11-4
 - cabling, Cat-5 ... 11-4
 - configuration ... 11-10
 - configuration planning ... 11-19
 - disabling a port ... 11-23
 - EPS, defined ... 11-3
 - event log messages ... 11-23
 - initial software release ... 11-2
 - insufficient power ... 11-6
 - maximum load calculation ... 11-21
 - messages ... 11-23
 - minimum power requirement ... 11-6, 11-7
 - MPS, defined ... 11-3
 - non-PoE device support ... 11-6
 - oversubscribed ... 11-3
 - PD, defined ... 11-3
 - pin pairs ... 11-4
 - port identifiers ... 11-12
 - power allocation ... 11-7
 - power supplies ... 11-2
 - priority class, defined ... 11-3
 - priority policies ... 11-20
 - priority, port ... 11-6, 11-8
 - PSE, defined ... 11-3
 - related publications ... 11-4
 - RPS, defined ... 11-3
 - security ... 11-20
 - setmib ... 11-12
 - status ... 15-52
 - supported switches ... 11-2
 - terminology ... 11-3
 - threshold, power ... 11-5, 11-11
 - unneeded power ... 11-6
 - viewing status ... 11-15
 - VLAN assignments ... 11-19
 - wire pairs, cable ... 11-4
- Port
 - duplex, view ... 10-8
 - speed, view ... 10-8
- port
 - address table ... B-13
 - auto negotiation ... 10-4
 - broadcast limit ... 10-14
 - CLI access ... 10-8
 - context level ... 10-10
 - counters ... B-10
 - counters, reset ... B-10
 - fiber-optic ... 10-5
 - MAC address ... D-4, D-5
 - menu access ... 10-6
 - monitoring ... B-23
 - monitoring, static LACP trunk ... B-23
 - monitoring, VLAN ... B-23
 - speed change, transceiver ... 10-4
 - traffic patterns ... B-10
 - trunk
 - See* port trunk.
 - utilization ... 5-17
 - web browser interface ... 5-17
 - web browser access ... 10-18
- Port Configuration ... 10-1, 13-1
- port names, friendly
 - configuring ... 10-19
 - displaying ... 10-21
 - summary ... 10-18
- port security
 - port trunk restriction ... 13-3
 - trunk restriction ... 13-8
- port trunk ... 13-2
 - See also* LACP.
 - bandwidth capacity ... 13-2
 - caution ... 13-3, 13-9, 13-17
 - CLI access ... 13-11
 - default trunk type ... 13-10
 - enabling dynamic LACP ... 13-15
 - IGMP ... 13-8
 - limit ... 13-2
 - limit, combined ... 13-19

- link requirements ... 13-3
- logical port ... 13-8
- media requirements ... 13-7
- media type ... 13-3
- menu access to static trunk ... 13-9
- monitor port restrictions ... 13-8
- monitoring ... B-23
- nonconsecutive ports ... 13-2
- port security restriction ... 13-8
- removing port from static trunk ... 13-15
- requirements ... 13-7
- SA/DA ... 13-25
- spanning tree protocol ... 13-8
- static trunk ... 13-7
- static trunk, overview ... 13-5
- static/dynamic limit ... 13-19
- STP ... 13-8
- STP operation ... 13-7
- traffic distribution ... 13-7
- Trk1 ... 13-7
- trunk (non-protocol) option ... 13-6
- trunk option described ... 13-24
- types ... 13-6
- VLAN ... 13-8
- VLAN operation ... 13-7
- web browser access ... 13-17
- port trunk group
 - interface access ... 13-1
- port, active ... 15-25
- port-based access control
 - event log ... C-15
 - LACP not allowed ... 13-21
 - rules of operation ... 12-14, 12-31
 - troubleshooting ... C-15
- power interruption, effect on event log ... C-27
- power-over-ethernet
 - See* PoE.
- Power-Sourcing Equipment ... 15-27
- Procurve
 - support URL ... 5-13
- prompt, => ... C-58
- PSAP ... 15-27
- PSE ... 15-27
- Public Safety Answering Point ... 15-27
- public SNMP community ... 15-5, 15-13

Q

- QoS resources ... 14-5
- quick configuration ... 3-8
- quick start ... 1-8, 8-3

R

- RADIUS, web browser access ... 5-8
- rate-limiting
 - caution ... 14-4
 - configuration ... 14-5, 14-14
 - displaying configuration ... 14-6, 14-14
 - edge ports ... 14-4
 - effect of flow control ... 14-7, 14-18
 - effect on port trunks ... 14-9
 - exceeding configured rate ... 14-5, 14-9, 14-13, 14-18
 - intended use ... 14-4
 - note on testing ... 14-8, 14-17
 - operating notes ... 14-7, 14-18
 - operation ... 14-5, 14-12
 - optimum packet size ... 14-9, 14-19
 - per-port only ... 14-4
 - purpose ... 14-4
 - resource consumption ... 14-5
 - security ... 14-5, 14-13
 - traffic filters ... 14-9, 14-19
- reboot ... 3-8, 3-10, 3-12
 - See also* boot.
- reboot time ... 6-18
- reboot, actions causing ... 6-4
- rebooting the switch
- reconfigure ... 3-10
- reload ... 6-4, 6-18, 6-20
- remote session, terminate ... 7-8
- reset ... 3-12, 6-11
- Reset button ... 6-4
 - restoring factory default configuration ... C-58
- reset port counters ... B-10
- resetting the switch
 - factory default reset ... C-57
- restricted write access ... 15-12
- RFC
 - See* MIB.
- RFC 1493 ... 15-4
- RFC 1515 ... 15-4
- RFC 2922 ... 15-31
- RFC2737 ... 15-31

- RFC2863 ... 15-31
- RIP
 - debug command ... C-37
- RIP broadcast traffic, broadcast traffic, RIP ... 10-5, 10-14
- RMON ... 15-4
- RMON groups supported ... 15-22
- router
 - gateway ... 8-6
- router, hop ... 8-10
- routing
 - gateway fails ... C-11
 - OSPF debug ... C-37
 - RIP debug ... C-37
 - traceroute ... C-53
- RS-232 ... 2-4
- running-config, viewing ... 6-6
 - See also* configuration.

S

- SCP/SFTP
 - session limit ... A-15
- secure copy
 - See* SCP/SFTP.
- secure FTP
 - See* SCP/SFTP.
- security ... 5-11, 7-3
 - username and password ... 5-8
 - web browser access, RADIUS ... 5-8
- Self Test LED
 - behavior during factory default reset ... C-58
- serial number ... B-6
- setmib ... 11-12
- setmib, delay interval ... 15-39
- setmib, reinit delay ... 15-41
- setting fault detection policy ... 5-24
- setup screen ... 1-8, 8-3
- severity code, event log ... C-27
- sftp/scp, multiple config file ... 6-39
- show management ... 8-7, 9-10, 9-19
- show tech ... C-50
- slow network ... C-7
- SNMP ... 15-3
 - CLI commands ... 15-12
 - communities ... 15-4, 15-5, 15-12, 15-13
 - Communities screen ... 15-11
 - configure ... 15-4, 15-5

- IP ... 15-3
- notification, LLDP
 - SNMP notification ... 15-29
- public community ... 15-5, 15-13
- setmib ... 11-12
- thresholds ... 15-19
- traps ... 15-4, 15-19
- traps, well-known ... 15-19
- walkmib ... D-5, D-6
- SNMP communities
 - configuring with the CLI ... 15-14
 - configuring with the menu ... 15-13
- SNMP trap, LLDP ... 15-41
- SNMPv3
 - "public" community access caution ... 15-6
 - access ... 15-5
 - assigning users to groups ... 15-7
 - communities ... 15-11
 - enable command ... 15-7
 - enabling ... 15-6
 - group access levels ... 15-10, 15-11
 - groups ... 15-9
 - network management problems with snmpv3
 - only ... 15-6
 - notification ... 15-16
 - restricted-access option ... 15-6
 - set up ... 15-5
 - traps ... 15-16
 - users ... 15-5
- SNTp
 - broadcast mode ... 9-3, 9-11
 - broadcast mode, requirement ... 9-3
 - configuration ... 9-5
 - disabling ... 9-12
 - enabling and disabling ... 9-10
 - event log messages ... 9-28
 - manual config priority ... 8-12
 - menu interface operation ... 9-28
 - operating modes ... 9-3
 - poll interval ... 9-13
 - See also* TimeP.
 - selecting ... 9-3
 - show management ... 9-10
 - unicast mode ... 9-3, 9-11
 - unicast time polling ... 9-25
 - unicast, address priority ... 9-25
 - unicast, deleting addresses ... 9-27
 - unicast, replacing servers ... 9-27

- viewing ... 9-5, 9-9
- software
 - See* switch software.
- software image
 - See* switch software.
- software version ... B-6
- sorting alert log entries ... 5-20
- source-port filter ... 14-33
- spanning tree
 - fast-uplink, troubleshooting ... C-20
 - global information ... B-17
 - information screen ... B-17
 - problems related to ... C-20
 - show tech, copy output ... C-51
 - statistics ... B-17
 - using with port trunking ... 13-8
- special ports used ... 12-7
- specific VID ... 12-19
- SSH
 - TACACS exclusion ... A-14
 - troubleshooting ... C-20
- stack management
 - See* stacking.
- stacking ... 3-5, 3-6, 3-12, 3-14
- standard MIB ... 15-4
- starting a console session ... 3-4
- startup-config, viewing ... 6-6
 - See also* configuration.
- statistics ... 3-7, B-4
- statistics, clear counters ... 3-12, 6-11
- status and counters
 - access from console ... 3-7
- status and counters menu ... B-5
- status overview screen ... 5-6
- subnet ... 8-3, 8-8
- subnet mask ... 8-4, 8-6
 - See also* IP.
- support
 - URL ... 5-12
 - URL Window ... 5-12
- switch console
 - See* console.
- switch setup menu ... 3-8
- switch software
 - download using TFTP ... A-4
 - download, failure indication ... A-22
 - download, switch-to-switch ... A-18
 - download, troubleshooting ... A-21

- download, using TFTP ... A-4
- software image ... A-3
- version ... A-6, A-17
- Syslog
 - configure server IP ... C-36
 - configuring messaging ... C-40
 - facility, user ... C-44
 - logging command ... C-36, C-38
 - operating notes ... C-44
 - operation ... C-39
 - See also* debug command.
 - servers ... C-34
 - severity, "debug" ... C-38, C-44
- system configuration screen ... 7-9
- System Name parameter ... 7-10

T

- TACACS
 - SSH exclusion ... A-14
- Telnet ... 3-4
 - terminate session, kill command ... 7-8
- Telnet, enable/disable ... 7-4
- Telnet, outbound ... 7-6
- Telnet, problem ... C-6
- terminal access, lose connectivity ... 7-6
- terminal type ... 7-3
- terminate remote session ... 7-8
- TFTP
 - download ... A-5
- threshold setting ... 15-5, 15-13
- thresholds, SNMP ... 15-19
- time format ... C-27
- time protocol
 - selecting ... 9-3
- time server ... 8-2
- time zone ... 7-10, 7-13
- time, configure ... 7-13
- TimeP ... 8-3, 8-5
 - assignment methods ... 9-2
 - disabling ... 9-23
 - enabling and disabling ... 9-20
 - manual config priority ... 8-12
 - poll interval ... 9-23
 - selecting ... 9-3
 - server address listing ... 9-10, 9-19
 - show management ... 9-19
 - viewing and configuring, menu ... 9-17

- viewing, CLI ... 9-19
- timesync, disabling ... 9-23
- Time-To-Live ... 8-3, 8-5, 8-6, 8-10
 - See also* TTL.
- time-to-live, LLDP ... 15-28
- Time-To-Live, on primary VLAN ... 8-4
- TLV ... 15-27
- TLVs, mandatory ... 15-71
- top talker ... 13-26
- traceroute
 - asterisk ... C-55
 - blocked route ... C-56
 - fails ... C-54
- traffic monitoring ... 15-5, 15-13
- traffic, monitoring ... B-23
- traffic, port ... B-10
- transceiver, fiber-optic ... 10-5
- transceiver, speed change ... 10-4
- trap ... 5-25
 - authentication ... 15-19
 - authentication trap ... 15-22
 - CLI access ... 15-19
 - event levels ... 15-21
 - limit ... 15-19
 - receiver ... 15-19
 - SNMP ... 15-19
- trap notification ... 15-41
- trap receiver ... 15-4, 15-5
 - configuring ... 15-19, 15-20
- troubleshooting
 - ACL ... C-8
 - approaches ... C-4
 - browsing the configuration file ... C-50
 - console access problems ... C-5
 - diagnosing unusual network activity ... C-7
 - diagnostics tools ... C-45
 - fast-uplink ... C-19
 - ping and link tests ... C-45
 - restoring factory default configuration ... C-57
 - spanning tree ... C-19
 - switch software download ... A-21
 - switch won't reboot, shows => prompt ... C-58
 - unusual network activity ... C-7
 - using the event log ... C-27
 - web browser access problems ... C-5
- troubleshooting, SSH. ... C-20
- trunk
 - See* port trunk.

- TTL ... 8-3, 8-5, 8-6, 8-7
 - manual config priority ... 8-12
 - on primary VLAN ... 8-4
 - See also* Time-To-Live.
- TTL, IP ... 8-10
- TTL, LLDP ... 15-28
- Type-Length-Value ... 15-27
- types of alert log entries ... 5-21

U

- unauthorized access ... 15-22
- undersize packets ... 14-34
- Universal Resource Locator
 - See* URL.
- Unix, Bootp ... 8-13
- unrestricted write access ... 15-12
- unusual network activity ... C-7
- up time ... B-6
- URL ... 15-4
 - browser interface online help location ... 5-13
 - management ... 5-13
 - management server ... 5-12, 5-13
 - Procurve ... 5-13, 15-4
 - support ... 5-12, 5-13
- user name, using for browser or console
 - access ... 5-8, 5-10
- users, SNMPv3
 - See* SNMPv3.
- using extended CLI ... 12-27
- using the passwords ... 5-10
- utilization, port ... 5-17

V

- version, OS ... A-19
- version, switch software ... A-6, A-17
- view
 - duplex ... 10-8
 - Port speed ... 10-8
- VLAN ... 8-4, C-25
 - address ... 15-3
 - Bootp ... 8-13
 - configuring Bootp ... 8-13
 - device not seen ... C-24
 - event log entries ... C-27
 - ID ... 4-14
 - link blocked ... C-20

- MAC address ... D-2, D-5
- management and jumbo packets ... 14-32
- management VLAN, SNMP block ... 15-3
- monitoring ... B-3, B-23
- multinet ... 8-3
- multinetting ... 8-3, 8-8
- multiple ... 15-3
- multiple IP addresses ... 8-3, 8-8
- port configuration ... C-24
- primary ... 8-3
- reboot required ... 3-8
- same MAC, multiple VLANs ... D-6
- subnet ... 8-3, 8-8
- support enable/disable ... 3-8
- switch software download ... A-4
- tagging broadcast, multicast, and unicast traffic ... C-24
- VLAN ID
 - See* VLAN.
- VoIP ... 15-47
- VT-100 terminal ... 7-3

W

- walknib ... 15-31, D-5, D-6
- warranty ... 1-ii
- web agent enabled ... 5-2
- web agent,
 - advantages ... 2-5
- web browser access configuration ... 7-3
- web browser enable/disable ... 7-4
- web browser interface ... 2-5
 - access parameters ... 5-8
 - alert log ... 5-6, 5-20
 - alert log details ... 5-21
 - bandwidth adjustment ... 5-18
 - bar graph adjustment ... 5-18
 - disable access ... 5-2
 - enabling ... 5-4
 - error packets ... 5-17
 - fault detection policy ... 5-8, 5-24
 - fault detection window ... 5-24
 - features ... 2-5
 - first-time install ... 5-7
 - first-time tasks ... 5-7
 - main screen ... 5-16
 - online help ... 5-13
 - online help location specifying ... 5-13

- online help, inoperable ... 5-13
- overview ... 5-16
- Overview window ... 5-16
- password lost ... 5-10
- password, setting ... 5-9
- port status ... 5-19
- port utilization ... 5-17
- port utilization and status displays ... 5-17
- screen elements ... 5-16
- security ... 5-2, 5-8
- standalone ... 5-4
- status bar ... 5-22
- status indicators ... 5-23
- status overview screen ... 5-6
- system requirements ... 5-4
- troubleshooting access problems ... C-5
- URL default ... 5-13
- URL, management server ... 5-14
- URL, support ... 5-14
- web site, HP ... 15-4
- world wide web site, ProCurve
 - See* Procurve.
- write access ... 15-12
- write memory, effect on menu interface ... 3-13

X

- Xmodem OS download ... A-16

—This page is intentionally unused—



Technical information in this document
is subject to change without notice.

© Copyright 2006
Hewlett-Packard Development Company, L.P.
Reproduction, adaptation, or translation
without prior written permission is prohibited
except as allowed under the copyright laws.

October 2006

Manual Part Number
5990-6050